Next-to-leading order QCD predictions for W+3-jet distributions at hadron colliders

Citation: Berger, C.F. et al "Next-to-leading order QCD predictions for W+3-jet distributions at hadron colliders." Physical Review D 80.7 (2009): 074036. © 2009 The American Physical Society

As Published: http://dx.doi.org/10.1103/PhysRevD.80.074036

Publisher: American Physical Society

Persistent URL: http://hdl.handle.net/1721.1/52697

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

PHYSICAL REVIEW D **80,** 074036 (2009)

Next-to-leading order QCD predictions for W + 3-jet distributions at hadron colliders

C. F. Berger,[1] Z. Bern,[2] L. J. Dixon,[3] F. Febres Cordero,[2] D. Forde,[3] T. Gleisberg,[3] H. Ita,[2] D. A. Kosower,[4] and D. Maître[5]

[1]*Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*
[2]*Department of Physics and Astronomy, UCLA, Los Angeles, California 90095-1547, USA*
[3]*SLAC National Accelerator Laboratory, Stanford University, Stanford, California 94309, USA*
[4]*Institut de Physique Théorique, CEA–Saclay, F–91191 Gif-sur-Yvette cedex, France*
[5]*Department of Physics, University of Durham, DH1 3LE, United Kingdom*

(Received 24 July 2009; published 29 October 2009)

We present next-to-leading order QCD predictions for a variety of distributions in W + 3-jet production at both the Tevatron and the Large Hadron Collider. We include all subprocesses and incorporate the decay of the W boson into leptons. Our results are in excellent agreement with existing Tevatron data and provide the first quantitatively precise next-to-leading order predictions for the LHC. We include all terms in an expansion in the number of colors, confirming that the specific leading-color approximation used in our previous study is accurate to within 3%. The dependence of the cross section on renormalization and factorization scales is reduced significantly with respect to a leading-order calculation. We study different dynamical scale choices, and find that the total transverse energy is significantly better than choices used in previous phenomenological studies. We compute the one-loop matrix elements using on-shell methods, as numerically implemented in the BLACKHAT code. The remaining parts of the calculation, including generation of the real-emission contributions and integration over phase space, are handled by the SHERPA package.

DOI: 10.1103/PhysRevD.80.074036 PACS numbers: 12.38.−t, 12.38.Bx, 13.87.−a, 14.70.−e

I. INTRODUCTION

The upcoming start of physics runs at the LHC has added impetus to the long-standing quest to improve theoretical control over standard model backgrounds to new physics searches at hadron colliders. Some backgrounds can be understood without much theoretical input. For example, a light Higgs boson decaying into two photons produces a narrow bump in the diphoton invariant mass, from which the large but smooth QCD background can be subtracted experimentally using sideband information. However, for many searches the signals are excesses in broader distributions of jets, along with missing energy and charged leptons or photons; such searches require a much more detailed theoretical understanding of the QCD backgrounds. A classic example is the production of a Higgs boson in association with a W boson at the Tevatron, with the Higgs decaying to a $b\bar{b}$ pair, and the W decaying to a charged lepton and a neutrino. The peak in the $b\bar{b}$ invariant mass is much broader than in the diphoton one; therefore variations in the backgrounds, including QCD production of $Wb\bar{b}$, across the signal region are more difficult to assess.

In this paper, we focus on a related important class of backgrounds, production of multiple (untagged) jets in association with a W boson. Such events, with a leptonically decaying W, form a background to supersymmetry searches at the LHC that require a lepton, missing transverse energy, and jets [1]. If the lepton is missed, they also contribute to the background for similar searches not requiring a lepton. The rate of events containing a W along with multiple jets can be used to calibrate the corresponding rate for Z production with multiple jets, which form another important source of missing transverse energy when the Z decays to a pair of neutrinos. Analysis of W plus multi-jet production will also assist in separating these events from the production of top-quark pairs, so that more detailed studies of the latter can be performed.

The first step toward a theoretical understanding of QCD backgrounds is the evaluation of the cross section at leading order (LO) in the strong coupling α_S. Our particular focus is on high jet multiplicity in association with vector boson production. Many computer codes [2–4] are available to generate predictions at leading order. Some of the codes incorporate higher-multiplicity leading-order matrix elements into parton showering programs [5,6], using matching (or merging) procedures [7,8]. LO predictions suffer from large renormalization- and factorization-scale dependence, growing with increasing jet multiplicity, and already up to a factor of 2 in the processes we shall study. Next-to-leading order (NLO) corrections are necessary to obtain quantitatively reliable predictions. They typically display a greatly reduced scale dependence [9]. Fixed-order results at NLO can also be matched to parton showers. This has been done for a growing list of processes within the MC@NLO program and the POWHEG method [10]. It would be desirable to extend this matching to higher-multiplicity processes such as those we discuss in the present paper.

The production of massive vector bosons in association with jets at hadron colliders has been the subject of theoretical studies for over three decades. Early studies were of large transverse-momentum muon-pair production at lead-

ing order [11], followed by the leading-order matrix elements for W + 2-jet production [12] and corresponding phenomenological studies [13,14]. The early leading-order studies were followed by NLO predictions for vector boson production in association with a single jet [15,16]. Leading-order results for vector-boson production accompanied by three or four jets appeared soon thereafter [17]. These theoretical studies played an important role in the discovery of the top quark [18]. Modern matrix element generators [2–4] allow for even larger numbers of final-state jets at LO. The one-loop matrix elements for W + 2-jet and Z + 2-jet production were determined [19] via the unitarity method [20] (see also Ref. [21]), and incorporated into the parton-level MCFM [22] code.

Studies of W production in association with heavy quarks have also been performed. NLO results for hadronic production of a W and a charm quark first appeared in Ref. [23]. More recently, NLO results have been presented for $W + b +$ jet production [24], as well as for $Wb\bar{b}$ production with full b quark mass effects [25]. The last two computations were combined to produce a full description of W production in association with a single b-jet in Ref. [26].

NLO studies of W production in association with more jets have long been desirable. However, a bottleneck to these studies was posed by one-loop amplitudes involving six or more partons [9]. On-shell methods [27], which exploit unitarity and recursion relations, have successfully broken this bottleneck, by avoiding gauge-noninvariant intermediate steps, and reducing the problem to much smaller elements analogous to tree-level amplitudes. Approaches based on Feynman diagrams have also led to new results with six external partons, exemplified by the NLO cross section for producing $t\bar{t}b\bar{b}$ at hadron colliders [28]. We expect that on-shell methods will be particularly advantageous for processes involving many external gluons, which often dominate multijet final states. Various results [29–33] already indicate the suitability of these methods for a general-purpose numerical approach to high-multiplicity one-loop amplitudes.

We recently presented the first NLO results for W + 3-jet production including all subprocesses [34], using one-loop amplitudes obtained by on-shell methods. This study used a specific type of leading-color approximation designed to have small corrections—under 3%, as verified in $W + 1, 2$-jet production—while reducing the required computer time. The study was performed for the Tevatron, with the same cuts employed by the CDF collaboration in their measurement of $W + n$-jet production [35]. The NLO corrections show a much-reduced dependence on the renormalization and factorization scales, and excellent agreement with the CDF data for the distribution in the transverse energy E_T of the third-most energetic jet.

In the present paper, we continue our study of W + 3-jet production. We present results for W + 3-jet production at the LHC as well as at the Tevatron. As before, we include all subprocesses and take all quarks to be massless. (We do not include top-quark contributions, but expect them to be very small for the distributions we shall present.) We extend the previous results by including specific virtual contributions that are subleading in the number of colors, which we had previously neglected. We shall demonstrate that, as expected, these subleading-color corrections to cross sections and distributions are uniformly small, generally under 3%. We present three sets of distributions at the Tevatron: the E_T of the third most energetic jet, the total transverse energy H_T [36], and the dijet invariant masses. These distributions are again computed with the same cuts used by CDF. (As discussed further in Sec. III, we used the infrared-safe SISCONE jet algorithm [37] in place of JETCLU, the cone algorithm used by CDF.) The code we use is general-purpose, permitting the analysis at NLO of any infrared-safe observable in W + 3-jet events. We also present a wide variety of distributions for the ultimate LHC energy of 14 TeV. We find that all the NLO cross sections and distributions display the expected reduction in renormalization- and factorization-scale dependence compared to the same quantities calculated at leading order.

The shapes of distributions at leading order are quite sensitive to the functional form of the scale choice. As expected, the change in shape between LO and NLO distributions can be reduced by choosing typical energy scales event-by-event for the renormalization and factorization scales, as noted by many authors over the years [14,16,38]. The vector boson transverse energy E_T^W, employed as an event-by-event scale in previous predictions and comparisons with data [14,34,35,39], turns out to be a poor characterization of the typical energy scale for events with large jet transverse energies, as at the LHC. We find that the total partonic transverse energy is a much better choice. Recently, similar deficiencies in the scale choice of E_T^W at LO have been observed independently, and another variable, related to the invariant mass of the final-state jets, has been proposed as a replacement [40]. Here we go further and demonstrate that for LHC energies, E_T^W is a poor scale choice not only at LO but also at NLO, yielding negative differential cross sections in the tails of some distributions. This pathology arises from large residual logarithms induced by disparities between momentum-transfer scales in multijet processes and the value of E_T^W.

For W + 3-jet production, choosing the total partonic transverse energy as the scale gives rise to shapes of distributions at LO that are typically similar to those at NLO. For a few W + 3-jet distributions genuine NLO effects are present, and significant shape changes remain between LO and NLO. These differences are usually less pronounced than in $W + 1, 2$-jet production. In the latter cases, the LO kinematics are more constrained, leading to significantly larger NLO corrections. In any event, an accurate description of the shape of any distribution re-

quires an NLO computation, either to confirm that its shape is unmodified compared to LO, or to quantitatively determine the shape change.

Ellis *et al.* have recently presented partial NLO results for $W + 3$-jet production. Their first calculation [41] was restricted to leading-color contributions to two-quark subprocesses, rendering it unsuitable for phenomenological studies. Their version of the leading-color approximation drops subleading-color terms in both the virtual and real-emission contributions. Quite recently [42] the same authors have added the leading-color contributions from four-quark processes, folding in the decay of the W in the zero-width approximation. They extended their leading-color approximation to include n_f-dependent terms, and estimated the full-color result based on the leading-order ratio of the full-color (FC) and leading-color (LC) cross sections. The value of the double ratio $(\sigma^{\mathrm{NLO,FC}}/\sigma^{\mathrm{NLO,LC}})/(\sigma^{\mathrm{LO,FC}}/\sigma^{\mathrm{LO,LC}})$ they use implicitly is quite sensitive to the inclusion of n_f terms, and as noted by the authors, sensitive to cancellations between the two-quark and four-quark contributions. It is nonetheless interesting that their estimate for the total cross section is within a few percent of both our earlier result [34] and the full-color one presented in this paper. It would be interesting to test their estimates for various distributions against the complete results presented here; we leave such a comparison to future work.

Next-to-leading order cross sections are built from several ingredients: virtual corrections, computed from the interference of tree-level and one-loop amplitudes; real-emission corrections; and a mechanism for isolating and integrating the infrared singularities in the latter. We evaluate the one-loop amplitudes needed for $W + 3$-jet production at NLO using the BLACKHAT library [30]. This library implements on-shell methods for one-loop amplitudes numerically. Related methods have been implemented in several other programs [29,31–33,43,44]. A numerical approach to amplitudes requires attention to numerical instabilities induced by round-off error. We have previously verified BLACKHAT's stability for one-loop six-, seven-, and eight-gluon amplitudes [30], and for leading-color amplitudes for a vector boson with up to five partons [45], using a flat distribution of phase-space points. In the present work, we confirm the stability of BLACKHAT-computed matrix elements for an ensemble of points chosen in the same way as in the actual numerical integration of the cross section.

The real-emission corrections to the LO process arise from tree-level amplitudes with one additional parton: an additional gluon, or a quark–antiquark pair replacing a gluon. To isolate and cancel the infrared divergences that arise in the integration of these terms, we use the Catani–Seymour dipole subtraction method [46], as implemented [47] in the program AMEGIC++ [4], itself part of the SHERPA framework [6]. (We also use AMEGIC++ for the required tree-level matrix elements.) Other automated implementations of the dipole subtraction method have been presented recently [48].

The smallness of the subleading-color corrections to the specific leading-color approximation employed in Ref. [34] allows us to use a "color-expansion sampling" approach [49]. In this approach the subleading-color terms, while more time-consuming per phase-space point, are sampled with lower statistics than the leading-color ones, and therefore do not impose an undue burden on the computer time required.

This paper is organized as follows. In Sec. II we summarize our calculational setup, and demonstrate the numerical stability of the one-loop matrix elements. In Sec. III we present results for the Tevatron, using the same experimental cuts as CDF. In Sec. IV we discuss scale choices, showing that the choice of W transverse energy typically used for Tevatron studies can lead to significant distortions in the shapes of distributions over the broader range of kinematics accessible at the LHC. We advocate instead the use of scale choices that more accurately reflect typical energy scales in the process, such as the total partonic transverse energy, or a fixed fraction of it. In Sec. V, we present a wide variety of distributions for the LHC. We highlight two particular topics in subsequent sections. Section VI examines properties of the leptons produced by W decay in $W + 3$ jet events. The different pseudorapidity distributions for electrons and positrons are presented. Then we show the ratios, between W^+ and W^-, of the transverse energy distributions for both the charged leptons and neutrinos. These two ratios have strikingly different behavior at large E_T, presumably due to the effects of W polarization. In Sec. VII we present results for the jet-emission probability, as a function of the pseudorapidity separation of the leading two jets. These results are relevant for searches for the Higgs boson in vector-boson fusion production. In Sec. VIII, we discuss the specific leading-color approximation used in our previous study, and our approach to computing the subleading-color terms. We give our conclusions in Sec. IX. Finally, in an appendix we give values of squared matrix elements at a selected point in phase space.

II. CALCULATIONAL SETUP

At NLO, the $W + 3$-jet computation can be divided into four distinct contributions:

(i) the leading-order contribution, requiring the tree-level $W + 5$-parton matrix elements;
(ii) the virtual contribution, requiring the one-loop $W + 5$-parton matrix elements (built from the interference of one-loop and tree amplitudes);
(iii) the subtracted real-emission contribution, requiring the tree-level $W + 6$-parton matrix elements, an approximation capturing their singular behavior, and

integration of the difference over the additional-emission phase space;

(iv) the integrated approximation (real-subtraction term), whose infrared-singular terms must cancel the infrared singularities in the virtual contribution.

Each of these contributions must be integrated over the final-state phase space, after imposing appropriate cuts, and convoluted with the initial-state parton distribution functions.

We evaluate these different contributions using a number of tools. We compute the virtual corrections using on-shell methods, implemented numerically in BLACKHAT, as outlined below. The subtraction term is built using Catani–Seymour dipoles [46] as implemented [47] in AMEGIC++ [4]. This matrix-element generator is part of the SHERPA package [6]. AMEGIC++ also provides our tree-level matrix elements. The phase-space integration is handled by SHERPA, using a multichannel approach [50]. The SHERPA framework makes it simple to include various experimental cuts on phase space, and to construct and analyze a wide variety of distributions. With this setup, it is straightforward to make NLO predictions for any infrared-safe physical observable. We refer the reader to Refs. [4,6,47] for descriptions of AMEGIC++, SHERPA, and the implementation of the Catani–Seymour dipole subtraction method.

A. Subprocesses

The W + 3-jet process, followed by leptonic W decay,

$$\left.\begin{array}{l}\text{Tevatron: } p\bar{p} \\ \text{LHC: } pp\end{array}\right\} \to W^{\pm} + 3\,\text{jets} \to e^{\pm}\overset{(-)}{\nu}_e + 3\,\text{jets}, \quad (2.1)$$

receives contributions from several partonic subprocesses. At leading order, and in the virtual NLO contributions, these subprocesses are all obtained from

$$q\bar{q}'Q\bar{Q}g \to W^{\pm} \to e^{\pm}\overset{(-)}{\nu}_e, \quad (2.2)$$

$$q\bar{q}'ggg \to W^{\pm} \to e^{\pm}\overset{(-)}{\nu}_e, \quad (2.3)$$

by crossing three of the partons into the final state. The W couples to the q-q' line. We include the decay of the vector boson ($W^{\pm}$) into a lepton pair at the amplitude level. The W can be off shell; the lepton-pair invariant mass is drawn from a relativistic Breit-Wigner distribution whose width is determined by the W decay rate Γ_W. For definiteness we present results for W bosons decaying to either electrons or positrons (plus neutrinos). We take the leptonic decay products to be massless; in this approximation the corresponding results for μ (and τ) final states are of course identical. Amplitudes containing identical quarks are generated by antisymmetrizing in the exchange of appropriate q and Q labels. The light quarks, u, d, c, s, b, are all treated as massless. We do not include contributions to the amplitudes from a real or virtual top quark; its omission should have a very small effect on the overall result. Except as noted below, we use the same setup for the results we report for W + 1, 2-jet production.

B. Color organization of virtual matrix elements

To compute the production of W + 3 jets at NLO, we need the one-loop amplitudes for the processes listed in Eqs. (2.2) and (2.3). Amplitudes in gauge theories are naturally decomposed into a sum over permutations of terms; each term is the product of a color factor and a color-independent kinematic factor called a partial or color-ordered amplitude. It is convenient to decompose the one-loop amplitudes further, into a set of *primitive* amplitudes [19,51]. These are the basic gauge-invariant building blocks of the amplitude, in which the ordering of all colored external legs is fixed, the direction of fermion lines within the loop is fixed, and n_f terms arising from fermion loops are separated out. In BLACKHAT, the primitive amplitudes are computed directly using the on-shell methods reviewed in the next subsection. The primitive amplitudes are then combined to obtain the partial amplitudes. The virtual contributions are assembled by summing over interferences of the one-loop partial amplitudes with their tree-level counterparts.

In organizing the amplitude, it is useful to keep the numbers of colors N_c and of flavors n_f as parameters, setting them to their standard model values only upon evaluation. Matrix elements, whether at tree level or at one loop, can be organized in an expansion dictated by the $N_c \to \infty$ limit. In this expansion, the standard "leading-color" contribution is the coefficient of the leading power of N_c, and "subleading-color" refers to terms that are suppressed by at least one power of either $1/N_c^2$, or n_f/N_c from virtual quark loops. (The expansion in either quantity terminates at finite order, so if all terms are kept, the result is exact in $1/N_c$.)

Only one primitive amplitude contributes at leading order in $1/N_c$ to each leading-color partial amplitude. Figure 1 shows sample "parent" color-ordered Feynman diagrams for the leading-color primitive amplitudes needed for W + 3-jet production. Other diagrams contributing to a given primitive amplitude have fewer propagators in the loop. They can be obtained from the diagrams shown by moving vertices off of the loop onto trees attached to the loop, or by using four-gluon vertices, while

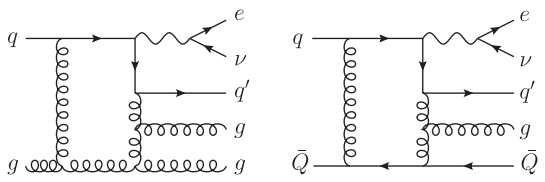

FIG. 1. Representative diagrams contributing at leading order in an expansion in the number of colors to the $qg \to e\nu q'gg$ and $q\bar{Q} \to e\nu q'g\bar{Q}$ one-loop amplitudes. The $e\nu$ pair couples to the quarks via a W boson.

preserving the cyclic (color) ordering of the external legs and the planar topology of the diagram. In the leading-color primitive amplitudes, the W boson is between the q and q' external legs, with no other partons in between.

In subleading-color terms, a greater number and variety of primitive amplitudes appear, and some primitive amplitudes contribute to more than one subleading-color partial amplitude. A few of the parent diagrams for subleading-color primitive amplitudes are shown in Fig. 2. In such diagrams, either another parton appears between the W boson and either q or q', or a gluon is emitted between Q and $\bar{Q}$ in process (2.2), or the diagram contains a closed fermion loop. In the present paper, we include *all* subleading-color contributions. In Sec. VIII, we discuss in greater detail how to evaluate the full virtual cross section efficiently by taking advantage of the smallness of the subleading-color contributions.

C. On-shell methods

The computation of one-loop partonic amplitudes has presented until recently a bottleneck to NLO predictions for hadronic production of four or more final-state objects (jets included). The on-shell method has broken this bottleneck. This approach is based on the unitarity method [20], including its newer refinements, together with on-shell recursion relations [52] at one loop [53]. The refinements [54–57] rely on the use of complex momenta, generalized unitarity and the analytic structure of integrands, as well as subtractions to make efficient use of the known basis of one-loop integrals. The one-loop matrix elements [19] used by the MCFM program [22] for NLO predictions of W + 2-jet production were computed analytically using an early version of this approach, and indeed served to introduce the use of generalized unitarity [58] as an efficient technique for loop computations. As applied to hadron colliders, these matrix elements have three final-state objects. Feynman-diagram calculations have also reached into this domain [9]. Beyond this, improved integral reduction techniques [59] have even made possible the computation of matrix elements for four final-state objects [28,60] and NLO predictions using them.

Nonetheless, textbook Feynman-diagrammatic approaches suffer from a factorial increase in complexity (or exponential if color ordered) and increasing degree of tensor integrals, with increasing number of external legs. The unitarity method for one-loop amplitudes, in contrast, can be cast in a form with only a polynomial increase in complexity per color-ordered helicity configuration [30,31,61]. This suggests that it will have an increasing advantage with increasing jet multiplicity. At fixed multiplicity, on-shell methods gain their improved efficiency by removing *ab initio* the cancellation of gauge-variant terms, and eliminating the need for tensor-integral (or higher-point integral) reductions. The problem is reduced to the computation of certain rational functions of the kinematic variables, to which efficient treelike techniques can be applied. On-shell methods have also led to a host of analytic results, including one-loop amplitudes in QCD with an arbitrary number of external legs, for particular helicity assignments [53,61]. The reader may find reviews and further references in Refs. [9,27,62].

The BLACKHAT library implements on-shell methods for one-loop amplitudes numerically. We have described the computation of amplitudes using BLACKHAT elsewhere [30,45]. We limit ourselves here to an overview, along with a discussion of new features that arise when we include subleading-color contributions to the cross section.

Any one-loop amplitude can be written as a sum of terms containing branch cuts in kinematic invariants, C_n, and terms free of branch cuts, R_n,

$$A_n = C_n + R_n. \tag{2.4}$$

The cut part C_n can in turn be written as a sum over a basis of scalar integrals [63],

$$C_n = \sum_i d_i I_4^i + \sum_i c_i I_3^i + \sum_i b_i I_2^i. \tag{2.5}$$

The scalar integrals $I^i_{2,3,4}$—bubbles, triangles, and boxes—are known functions [64]. They contain all the amplitude's branch cuts, packaged inside logarithms and dilogarithms. (Massive particles propagating in the loop would require the addition of tadpole contributions.) We take all external momenta to be four dimensional. Following the spinor-helicity method [12,65], we can then reexpress all external momenta in terms of spinors. The coefficients of these integrals, b_i, c_i, and d_i, as well as the rational remainder R_n, are then all rational functions of spinor variables, and more specifically of spinor products. The problem of cal-

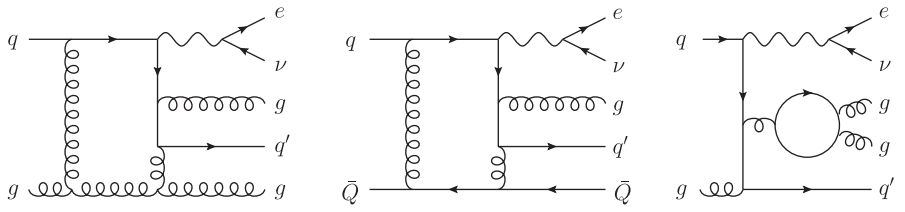

FIG. 2. Representative diagrams contributing only at subleading order in an expansion in the number of colors to the $qg \to e\nu q' gg$ and $q\bar{Q} \to e\nu q' g\bar{Q}$ one-loop amplitudes. In such contributions, either an external gluon, or a gluon splitting to a $\bar{Q}Q$ pair, is emitted from the q-q' line, between the W boson and one of the external quarks, q or q', in the cyclic ordering of the external legs.

culating a one-loop amplitude then reduces to the problem of determining these rational functions.

Generalized unitarity improves upon the original unitarity approach by isolating smaller sets of terms, hence making use of simpler on-shell amplitudes as basic building blocks. Furthermore, by isolating different integrals, it removes the need for integral reductions; and by computing the coefficients of scalar integrals directly, it removes the need for tensor reductions. Britto, Cachazo and Feng [54] showed how to combine generalized unitarity with a twistor-inspired [66] use of complex momenta to express all box coefficients as a simple sum of products of tree amplitudes. Forde [57] showed how to extend the technique to triangle and bubble coefficients. His method uses a complex parametrization and isolates the coefficients at specific universal poles in the complex plane. It is well suited to analytic calculation. Upon trading series expansion at infinity for exact contour integration via discrete Fourier summation [30], the method can be applied to numerical calculation as well, where it is intrinsically stable. Generalized unitarity also meshes well with the subtraction approach to integral reduction introduced by Ossola, Papadopoulos, and Pittau (OPP) [55]. As described in Ref. [30], in BLACKHAT we use Forde's analytic method, adapted to a numerical approach. We evaluate the boxes first, then the triangles, followed by the bubbles; the rational terms are computed separately. For each term computed by cuts, we enhance the numerical stability of Forde's method by subtracting prior cuts. This is similar in spirit to, though different in details from, one aspect of the OPP approach, in which all prior integral coefficients are subtracted at each stage.

The terms R_n, which are purely rational in the spinor variables, cannot be computed using four-dimensional unitarity methods. At present, there are two main choices for computing these contributions within a process-nonspecific numerical program: on-shell recursion and D-dimensional unitarity. Loop level on-shell recursion [53,61] is based on the tree-level on-shell recursion of Britto, Cachazo, Feng and Witten [52]. The utility of D-dimensional unitarity [43,44,62,67–69] grows out of the original observation [70] by van Neerven that dispersion integrals in dimensional regularization have no subtraction ambiguities. Accordingly the unitarity method in D dimensions retains all rational contributions to amplitudes [67]. This version of unitarity, in which tree amplitudes are evaluated in D dimensions, has been used in various analytic [68] and numerical [29,31–33,43,44,69,71] studies. We have implemented on-shell recursion in BLACKHAT, along with a ''massive continuation'' approach—related to D-dimensional unitarity—along the lines of Badger's method [72]. We speed up the on-shell recursion by explicitly evaluating some spurious poles analytically. Both approaches are used for our evaluation of the W + 3-jet virtual matrix elements. For producing the plots in this paper, we use on-shell recursion for the computation of primitive amplitudes with all negative helicities adjacent. These amplitudes have a simple pattern of spurious poles [61] (poles which cancel between the cut part C_n and rational part R_n). For them, on-shell recursion is faster than massive continuation in the present implementation.

BLACKHAT's use of four-dimensional momenta allows it to rely on powerful four-dimensional spinor techniques [12,65,73] to express the solutions for the loop momenta in generalized unitarity cuts in a numerically stable form [30]. In the computation of the rational terms using on-shell recursion, it also allows convenient choices for the complex momentum shifts. In four dimensions one can also employ simple forms of the tree amplitudes that serve as basic building blocks. While spinor methods arise most naturally in amplitudes with massless momenta, it is straightforward to include uncolored massive external states such as the W boson [12]. The methods are in fact quite general, and can also be applied usefully to one-loop amplitudes with internal massive particles, or external massive ones such as top quarks (treated in the narrow-width approximation) [43,74].

With the current version of BLACKHAT, the evaluation of a complete helicity-summed leading-color virtual interference term for a two-quark partonic subprocess (2.3), built out of all the primitive amplitudes, takes 530 ms on average for each phase-space point, on a 2.33 GHz Xeon processor. The evaluation of a complete four-quark partonic subprocess (2.2) with distinct quarks takes 185 ms (identical quarks take twice as long). The mix of subprocesses leads to an evaluation time of 470 ms on average for each phase-space point. (As described in Sec. II F, in performing the phase-space integration we sample a single subprocess at each point.) Using the ''color-expansion sampling'' approach we shall discuss in Sec. VIII, evaluating the subleading-color contributions would multiply this time by about 2.4, giving an average evaluation time of 1.1 s for the full color calculation.

D. Numerical stability of virtual contributions

BLACKHAT computes matrix elements numerically using on-shell methods. In certain regions of phase space, particularly near the vanishing loci of Gram determinants associated with the scalar integrals $I^i_{2,3,4}$, there can be large cancellations between different terms in the expansion (2.5) of the cut part C_n, or between the cut part and the rational part R_n in Eq. (2.4). There can also be numerical instabilities in individual terms. For example, the recursive evaluation of R_n includes a contribution from residues at spurious poles in the complex plane. These residues are computed by sampling points near the pole, in an approximation to a contour integral which can be spoiled if another pole is nearby.

In normal operation, BLACKHAT performs a series of tests to detect any unacceptable loss of precision. Whenever BLACKHAT detects such a loss, it reevaluates the problematic contributions to the amplitude (and only those terms) at that phase-space point using higher-precision arithmetic (performed by the QD package [75]). This approach avoids the need to analyze in detail the precise origin of instabilities and to devise work-arounds for each case. It does of course require that results be sufficiently stable, so that the use of higher precision is infrequent enough to incur only a modest increase in the overall evaluation time; this is indeed the case.

The simplest test of stability is checking whether the known infrared singularity of a given matrix element has been reproduced correctly. As explained in Ref. [30], this check can be extended naturally to check individual spurious-pole residues. Another test checks the accuracy of the vanishing of certain higher-rank tensor coefficients. From the interaction terms in the (renormalizable) QCD Lagrangian we know on general grounds which high-rank tensor coefficients have to vanish. All tensors with rank greater than m must vanish, for the m-point integrals with $m = 2, 3, 4$. If the integral corresponds to a cut line that is fermionic, then the maximum rank is reduced by one. In our approach the values of the higher-rank tensor coefficients may be computed without much extra cost in computation time. For a given generalized unitarity cut, when using complex loop-momentum parametrizations along the lines of Ref. [30], these tensor coefficients appear as coefficients of specific monomials in the complex parameters. Their values may be extracted as a by-product of evaluating the scalar integral coefficients. Similarly, in the massive continuation approach to computing the rational terms, particular tensor coefficients can be associated with specific monomials in the complex parameters and in an auxiliary complex mass parameter entering the loop-momentum parametrization.

We apply the latter check when computing coefficients of scalar bubble integrals, as well as bubble contributions to the rational terms in the massive continuation approach. The value of this check is twofold. First, it focuses on a small part of the computation, namely, single bubble coefficients. This allows BLACKHAT to recompute at higher precision just the numerically-unstable contributions, instead of the entire amplitude. By contrast, the above-mentioned check of the infrared singularity assesses the precision of the entire cut part C_n of the given primitive amplitude, and so it requires more recomputation if it fails. Second, the check applied to the bubble contributions in the massive continuation approach assesses the precision of the rational part R_n, which is inaccessible to the infrared-singularity check.

Finally, a further class of tests of numerical precision looks for large cancellations between different parts of A_n, in particular, between C_n and R_n in Eq. (2.4).

We have previously assessed the numerical stability of BLACKHAT for six-, seven-, and eight-gluon one-loop amplitudes [30], as well as for the leading-color amplitudes for a vector boson with up to five partons [45] used in the present study. These earlier studies used a flat phase-space distribution. Here we show the stability of BLACKHAT over phase-space points selected in the same way as in the computation of cross sections and distributions. As will be discussed in Sec. II F, the phase-space points are selected using an integration grid that has been adapted to the leading-order cross section.

In Fig. 3, we illustrate the numerical stability of the leading-color virtual interference term (or squared matrix element), $d\sigma_V$, summed over colors and over all helicity configurations for two subprocesses, $gd \to e^- \bar{\nu} d \bar{d} u$ and $gd \to e^- \bar{\nu} g g u$. (The grid here has been adapted to each of the subprocesses individually, instead of to the sum over subprocesses.) We have checked that the other subprocesses are similarly stable. The horizontal axis of Fig. 3 shows the logarithmic error,

$$\log_{10}\left(\frac{|d\sigma_V^{\text{num}} - d\sigma_V^{\text{target}}|}{|d\sigma_V^{\text{target}}|}\right), \tag{2.6}$$

for each of the three components: $1/\epsilon^2$, $1/\epsilon^1$ and ϵ^0, where $\epsilon = (4 - D)/2$ is the dimensional regularization parameter. The targets have been computed by BLACKHAT using multiprecision arithmetic with at least 32 decimal digits, and 64 if the point is deemed unstable. The overwhelming majority (99.9%) of events are computed to better than one part in 10^4—that is, to the left of the "-4" mark on the horizontal axis.

We have also examined distributions in which each bin is weighted by the requisite squared matrix element and Jacobian factors. We find that they have quite similar shapes to the unweighted distributions shown in Fig. 3. This implies that the few events with a relative error larger than 10^{-4} make only a small contribution to the total cross section. We have verified that the difference between normal and high-precision evaluation in the total cross section, as well as bin-by-bin for all distributions studied, is at least 3 orders of magnitude smaller than the corresponding numerical integration error.

E. Real-emission corrections

In addition to the virtual corrections to the cross section provided by BLACKHAT, a NLO calculation also requires the real-emission corrections to the LO process. These terms arise from tree-level amplitudes with one additional parton: an additional gluon, or a quark–antiquark pair replacing a gluon. Representative real-emission diagrams are shown in Fig. 4. Infrared singularities develop when the extra parton momentum is integrated over phase-space regions unresolved by the jet algorithm or jet cuts. The resulting singular integrals cancel against singular terms in the virtual corrections, and against counter-terms associ-

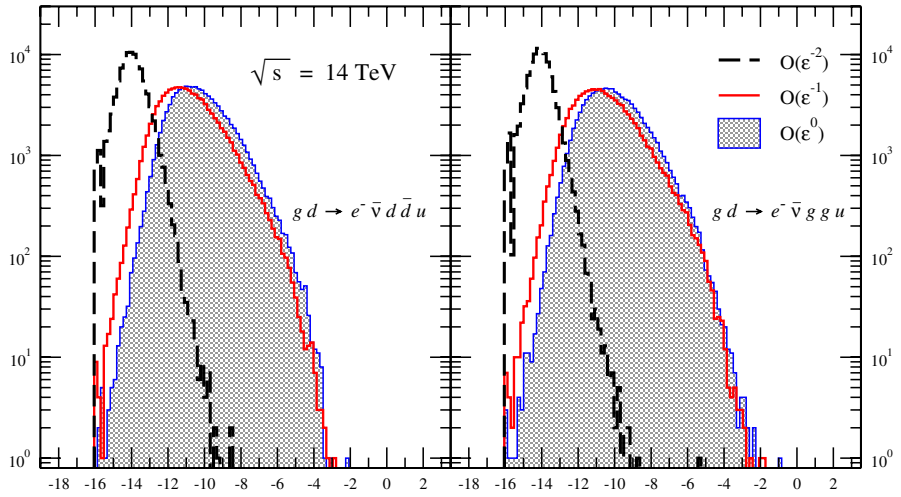

FIG. 3 (color online). The distribution of the relative error in the leading-color virtual cross section for two subprocesses, $gd \to e^- \bar{\nu} d \bar{d} u$ and $gd \to e^- \bar{\nu} g g u$. The phase-space points are selected in the same way as those used to compute cross sections at the LHC. The horizontal axis is the logarithm of the relative error (2.6) between an evaluation by BLACKHAT, running in production mode, and a target expression evaluated using higher precision with at least 32 decimal digits (or up to 64 decimal digits for unstable points). The vertical axis shows the number of events out of 100 000 with the corresponding error. The dashed (black) line shows the $1/\epsilon^2$ term; the solid (red) curve, the $1/\epsilon$ term; and the shaded (blue) curve, the finite (ϵ^0) term.

ated with the evolution of parton distributions. As mentioned above, to carry out these cancellations, we use the Catani–Seymour dipole subtraction method [46] as implemented [47] in the program AMEGIC++ [4], which is part of the SHERPA framework [6]. This implementation of dipole subtraction has already been tested [47] in explicit comparisons against the DISENT program [76].

The implementation introduces two free parameters, α_{cut} and α_{dipole}. The first, α_{cut}, parametrizes the volume of phase space to be cut out around the soft or collinear singularity. From an analytic point of view, α_{cut} could be taken to zero, as the cancellation of counterterms against the matrix element's singularities is exact. In numerical implementations, however, round-off error can spoil this cancellation. Previous studies have shown that the final result is independent of this cut-off parameter once it is sufficiently small [47]. We use $\alpha_{\text{cut}} = 10^{-8}$.

The second parameter, α_{dipole}, characterizes a common modification of subtraction terms in phase space *away* from the singularity [77], restricting the support of a given subtraction term to the vicinity of its singularity. This allows the program to compute only a subset of dipole terms, as many will now be identically zero at a given phase-space point. Because the number of dipole terms is large (scaling as m^3 for processes containing m partons), this reduces the computational burden considerably. We run our code with several different values of α_{dipole}, and check the independence of the final result on the value of α_{dipole} ($0 < \alpha_{\text{dipole}} \leq 1$). For example, the LHC W^- results for $\alpha_{\text{dipole}} = 0.03$ agree with those for $\alpha_{\text{dipole}} = 0.01$ to better than half a percent, within the integration errors. We have also run a large number of other lower-statistics checks demonstrating that cross sections are independent of the choice of α_{dipole}. Our default choice for the LHC is $\alpha_{\text{dipole}} = 0.03$, while for the Tevatron it is $\alpha_{\text{dipole}} = 0.01$.

F. Phase-space integration

Along with the automated generation of matrix elements and dipole terms, SHERPA also provides Monte Carlo integration methods. The phase-space generator combines *a priori* knowledge about the behavior of the integrands in phase space with self-adaptive integration methods. It employs a multichannel method in the spirit of Ref. [50]. Single channels (phase-space parametrizations) are generated by AMEGIC++ together with the tree-level matrix ele-

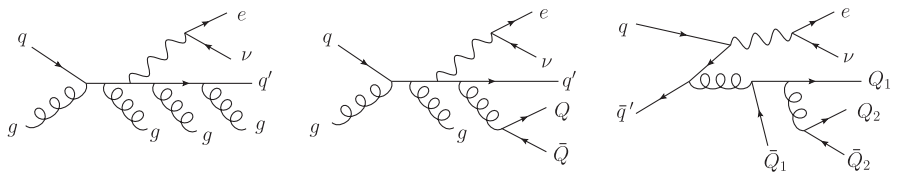

FIG. 4. Representative diagrams for the eight-point tree-level amplitudes, $qg \to e\nu q' ggg$, $qg \to e\nu q' g Q \bar{Q}$, and $q\bar{q}' \to e\nu Q_1 \bar{Q}_1 Q_2 \bar{Q}_2$. The $e\nu$ pair couples to the quarks via a W boson.

ments. Each parametrization reflects the structure of a Feynman amplitude, roughly reproducing its resonances, decay kinematics, and its soft and collinear structure. The most important phase-space parametrizations, determined by the adapted relative weight within the multichannel setup, are further refined using VEGAS [78].

The phase-space optimization (adaptation of channel weights and VEGAS grids) is performed in independent runs before the actual computation starts. The optimization is done on the sum of all contributing parton-level processes. We refer collectively to all the parameters of the optimization as the integration grid. Separate integration grids are constructed for the LO terms and for the real-emission contributions. To integrate the virtual contributions, we reuse the grid constructed for the LO terms. This procedure avoids the computational expense of evaluating the virtual terms merely for grid construction. The virtual to LO ratio is sufficiently flat across phase space that this results in only a slight inefficiency when evaluating distributions.

Following the initialization phase, the integration grids are frozen. In the ensuing production phase, we sample over subprocesses so that only a single parton-level subprocess is evaluated per phase-space point, selected with a probability proportional to its contribution to the total cross section. We choose to integrate the real-emission terms over about 10^8 phase-space points, the leading-color virtual parts over 2×10^6 phase-space points and the subleading-color virtual parts over 10^5 phase-space points. The LO and real-subtraction pieces are run separately with 10^7 points each. These numbers are chosen to achieve a total integration error of half a percent or less. For a given choice of scale μ, they give comparable running times for the real-emission and virtual contributions. Running times for leading- and subleading-color virtual contributions are also comparable.

G. Couplings and parton distributions

We work to leading order in the electroweak coupling and approximate the Cabibbo-Kobayashi-Maskawa (CKM) matrix by the unit matrix. This approximation causes a rather small change in total cross sections for the cuts we impose, as estimated by LO evaluations using the full CKM matrix. At the Tevatron, the full CKM results are about 1% smaller than with the unit CKM matrix; the difference is even smaller at the LHC. We express the W-boson couplings to fermions using the standard model input parameters shown in Table I. The parameter g_w^2 is derived from the others via,

$$g_w^2 = \frac{4\pi\alpha_{\rm QED}(M_Z)}{\sin^2\theta_W}. \tag{2.7}$$

TABLE I. Electroweak parameters used in this work.

parameter	value
$\alpha_{\rm QED}(M_Z)$	1/128.802
M_W	80.419 GeV
$\sin^2\theta_W$	0.230
Γ_W	2.06 GeV
g_w^2	0.4242 (calculated)

We use the CTEQ6M [79] parton distribution functions (PDFs) at NLO and the CTEQ6L1 set at LO. The value of the strong coupling is fixed accordingly, such that $\alpha_S(M_Z) = 0.118$ and $\alpha_S(M_Z) = 0.130$ at NLO and LO, respectively. We evolve $\alpha_S(\mu)$ using the QCD beta function for five massless quark flavors for $\mu < m_t$, and six flavors for $\mu > m_t$. (The CTEQ6 PDFs use a five-flavor scheme for all $\mu > m_b$, but we use the SHERPA default of six-flavor running above top-quark mass; the effect on the cross section is very small, on the order of 1% at larger scales.) At NLO we use two-loop running, and at LO, one-loop running.

H. Kinematics and observables

As our calculation is a parton-level one, we do not apply corrections due to nonperturbative effects such as those induced by the underlying event or hadronization. CDF has studied [35] these corrections at the Tevatron, and found they are under 10% when the $n^{\rm th}$ jet E_T is below 50 GeV, and under 5% at higher E_T.

For completeness we state the definitions of standard kinematic variables used to characterize scattering events. We denote the angular separation of two objects (partons, jets or leptons) by

$$\Delta R = \sqrt{(\Delta\phi)^2 + (\Delta\eta)^2}, \tag{2.8}$$

with $\Delta\phi$ the difference in the azimuthal angles, and $\Delta\eta$ the difference in the pseudorapidities. The pseudorapidity η is given by

$$\eta = -\ln\left(\tan\frac{\theta}{2}\right), \tag{2.9}$$

where θ is the polar angle with respect to the beam axis.

The transverse energies of massless outgoing partons and leptons, $E_T = \sqrt{p_x^2 + p_y^2}$, can be summed to give the total partonic transverse energy, $\hat{H}_T$, of the scattering process,

$$\hat{H}_T = \sum_p E_T^p + E_T^e + E_T^\nu. \tag{2.10}$$

All partons p and leptons are included in $\hat{H}_T$, whether or not they are inside jets that pass the cuts. We shall see in later sections that the variable $\hat{H}_T$ represents a good choice for the renormalization and factorization scale of a given event. Although the partonic version is not directly measurable, for practical purposes as a scale choice, it is

essentially equivalent (and identical at LO) to the more usual jet-based total transverse energy,

$$H_T = \sum_j E_{T,j}^{\rm jet} + E_T^e + E_T^\nu. \tag{2.11}$$

The partonic version $\hat{H}_T$ has the advantage that it is independent of the cuts; thus, loosening the cuts will not affect the value of the matrix element, because a renormalization scale of $\hat{H}_T$ will be unaffected. On the other hand, we use the jet-based quantity H_T, which is defined to include only jets passing all cuts, to compute observable distributions. Note that for $W + n$-jet production at LO, exactly n jets contribute to Eq. (2.11); at NLO either n or $n + 1$ jets may contribute.

The jet four-momenta are computed by summing the four-momenta of all partons that are clustered into them,

$$p_\mu^{\rm jet} = \sum_{i \in {\rm jet}} p_{i\mu}. \tag{2.12}$$

The transverse energy is then defined in the usual way, as the energy multiplied by the momentum unit vector projected onto the transverse plane,

$$E_T^{\rm jet} = E^{\rm jet} \sin\theta^{\rm jet}. \tag{2.13}$$

The total transverse energy as defined in Eq. (2.11) is intended to match the experimental quantity, given by the sum,

$$H_T^{\rm exp} = \sum_j E_{T,j}^{\rm jet} + E_T^e + \not{E}_T, \tag{2.14}$$

where $\not{E}_T$ is the missing transverse energy. Jet invariant masses are defined by

$$M_{ij}^2 = (p_i^{\rm jet} + p_j^{\rm jet})^2, \tag{2.15}$$

and the jets are always labeled $i, j = 1, 2, 3, \ldots$ in order of decreasing transverse energy E_T, with 1 being the leading (hardest) jet. The transverse mass of the W-boson is computed from the kinematics of its decay products, $W \to e\nu_e$,

$$M_T^W = \sqrt{2E_T^e E_T^\nu (1 - \cos(\Delta\phi_{e\nu}))}. \tag{2.16}$$

I. Checks

We have carried out numerous checks on our code, ranging from checks of the basic primitive amplitudes in specific regions of phase space to overall checks of total and differential distributions against existing codes. We have compared our results for the total cross section for $W + 1,2$-jet production (at a fixed scale $\mu = M_W$) with the results obtained from running MCFM [22]. Because the publicly available version of MCFM does not allow a cut in M_T^W we eliminated this cut in the comparison. (We had previously compared the matrix elements used in the latter code obtained from Ref. [19], to the results produced purely numerically in BLACKHAT.) Agreement at LO and NLO for $W + 1$-jet production at the LHC is good to a per mille level. For $W + 2$-jet production, at LO we find agreement with MCFM within a tenth of a percent, while at NLO, where the numerical integration is more difficult, we find agreement to better than half a percent.[1] We find the same level of agreement at NLO at the Tevatron, using a different set of cuts.[2]

We have carried out extensive validations of our code at a finer-grained level. We have confirmed that the code reproduces the expected infrared singularities (poles in ϵ) for the primitive amplitudes and the full color-dressed one-loop amplitudes [46,80]. We have also confirmed that the poles in ϵ in the full virtual cross section cancel against those found in the integrated real-subtraction terms [47].

We checked various factorization limits, both two-particle (collinear) and multiparticle poles. These factorization checks are natural in the context of on-shell recursion. This method constructs the rational terms using a subset of the collinear and multiparticle factorization poles; the behavior in other channels constitutes an independent cross check. For the leading-color primitive amplitudes, we verified that all factorization limits of the amplitudes are correct. (We also checked that all spurious poles cancel.) For the subleading-color primitive amplitudes, we verified the correct behavior as any two parton momenta become collinear. We also checked at least one collinear limit for each partial amplitude.

We had previously computed the leading-color amplitudes for the subprocess (2.3) in Ref. [45]. Ellis *et al.* [71] confirmed these values, and also computed the subleading-color primitive amplitudes. This evaluation used D-dimensional generalized unitarity [44,67,68], a decomposition of the processes in Eqs. (2.2) and (2.3) into primitive amplitudes [19,51], and the OPP formalism for obtaining coefficients of basis integrals [55]. We have compared the subleading-color primitive amplitudes at a selected phase-space point to the numerical values reported in Ref. [71], and find agreement, up to convention-dependent overall phases. Van Hameren, Papadopoulos, and Pittau (HPP) recently computed [33] the full helicity- and color-summed virtual cross section for the subprocess $u\bar{d} \to W^+ ggg$ at another phase-space point, for an undecayed on-shell W boson and including (small) virtual top-quark contributions. They used the OPP formalism and the CUTTOOLS [29] and HELAC-1L [3,33] codes. We have com-

[1]This level of agreement holds only for the most recent MCFM code, version 5.5. We thank John Campbell and Keith Ellis for assistance with this comparison. Here we matched MCFM by including approximate top-quark loop contributions, as given in Ref. [19], and we adopted MCFM's electroweak parameter conventions.

[2]In performing this comparison, we used a previous version of MCFM. The differences between the two versions at the Tevatron should be minor.

pared the full squared matrix element to the result produced by the HPP code, with the top-quark contributions removed.[3] We find agreement with their value of the ratio of this quantity to the LO cross section.[4] We have also found agreement with matrix-element results from the same code, allowing the W boson to decay to leptons, as in our setup. We give numerical values of the squared matrix elements for an independent set of subprocesses, evaluated at a different phase-space point, in an appendix.

As mentioned earlier, we verified that the computed values of the virtual terms are numerically stable when integrated over grids similar to those used for computing the cross section and distributions. We also checked that our integrated results do not depend on $\alpha_{\rm dipole}$, the unphysical parameter controlling the dipole subtraction [77], within integration uncertainties.

III. TEVATRON RESULTS

In this section we present next-to-leading order results for $W + 3$-jet production in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV, the experimental configuration at the Tevatron. We decay the W bosons into electrons or positrons (plus neutrinos) in order to match the CDF study [35]. In our earlier paper [34], we presented results for the third jet's transverse energy (E_T) distribution as well as the total transverse energy (H_T) distribution. Those calculations employed a particular leading-color approximation for the virtual terms [34]. As discussed in Sec. VIII, this approximation is an excellent one, accurate to within 3%. In the present paper, we give complete NLO results for a larger selection of distributions, including all subleading-color terms. It would be interesting to compare the new distributions with experimental results from both CDF and D0, as they become available.

We use the same jet cuts as in the CDF analysis [35],

$$E_T^{\rm jet} > 20 \text{ GeV}, \qquad |\eta^{\rm jet}| < 2. \tag{3.1}$$

Following Ref. [35], we quote total cross sections using a tighter jet cut, $E_T^{\rm jet} > 25$ GeV. We order jets by E_T. Both electron and positron final states are counted, using the same lepton cuts as CDF,

$$\begin{aligned} &E_T^e > 20 \text{ GeV}, \qquad |\eta^e| < 1.1, \\ &\not{E}_T > 30 \text{ GeV}, \qquad M_T^W > 20 \text{ GeV}. \end{aligned} \tag{3.2}$$

(We replace the $\not{E}_T$ cut by one on the neutrino E_T^ν.) CDF also imposes a minimum ΔR between the charged decay lepton and any jet; the effect of this cut, however, is undone by a specific acceptance correction [81]. Accordingly, we do not impose it.

For the LO and NLO results for the Tevatron we use an event-by-event common renormalization and factorization scale, set equal to the W boson transverse energy,

$$\mu = E_T^W \equiv \sqrt{M_W^2 + p_T^2(W)}. \tag{3.3}$$

To estimate the scale dependence we choose five values: $\mu/2$, $\mu/\sqrt{2}$, μ, $\sqrt{2}\mu$, 2μ.

The CDF analysis used the JETCLU cone algorithm [82] with cone radius $R = 0.4$. This algorithm is not generally infrared safe at NLO, so we use the seedless cone algorithm SISCONE [37] instead. Like other cone-type algorithms, SISCONE gives rise to jet-production cross sections that can depend on an overlap threshold or merging parameter, here called f. No dependence on f can develop at LO, because such dependence would require the presence of partons in the overlap of two cones. The $W + 1$-jet production cross section likewise cannot depend on f at NLO. We set this parameter to 0.5. (Unless stated otherwise we take this algorithm and parameter choice as our default.)

We expect similar results at the partonic level from any infrared-safe cone algorithm. For $W + 1, 2$-jet production we have confirmed that distributions using SISCONE are within a few percent of those obtained with MCFM using the midpoint cone algorithm [83]. (The midpoint algorithm is infrared-safe at NLO for $W + 1, 2$-jet production, but not for $W + 3$-jet production [37].) The algorithm dependence of $W + 3$-jet production at the Tevatron at NLO has also been discussed recently by Ellis *et al.* [42].

In Table II, we collect the results for the total cross section, comparing CDF data to the LO and NLO theoretical predictions computed using BLACKHAT and SHERPA. In both cases these are parton-level cross sections. Results from more sophisticated ("enhanced") LO analyses incorporating parton showering and matching schemes [7,8,84] may be found in Ref. [35]; however, large scale depen-

TABLE II. Total inclusive cross sections, in pb, for $W + n$ jets produced at the Tevatron with $W \to e\nu$ and $E_T^{n\text{th-jet}} > 25$ GeV, using the experimental cuts of Ref. [35]. The first column gives the experimental results as measured by CDF. The experimental statistical, systematic and luminosity uncertainties have been combined in quadrature. The second column shows LO results, and the third column the complete NLO results. In each case, the scale dependence is quoted in super- and subscripts and the numerical integration uncertainties in parentheses.

number of jets	CDF	LO	NLO
1	53.5 ± 5.6	$41.40(0.02)^{+7.59}_{-5.94}$	$57.83(0.12)^{+4.36}_{-4.00}$
2	6.8 ± 1.1	$6.159(0.004)^{+2.41}_{-1.58}$	$7.62(0.04)^{+0.62}_{-0.86}$
3	0.84 ± 0.24	$0.796(0.001)^{+0.488}_{-0.276}$	$0.882(0.005)^{+0.057}_{-0.138}$

[3]We thank Costas Papadopoulos and Roberto Pittau for providing us with these numbers.

[4]We can recover an undecayed W by integrating over the lepton phase space; that integral in turn can be done to high precision by replacing it with a discrete sum over carefully chosen points.

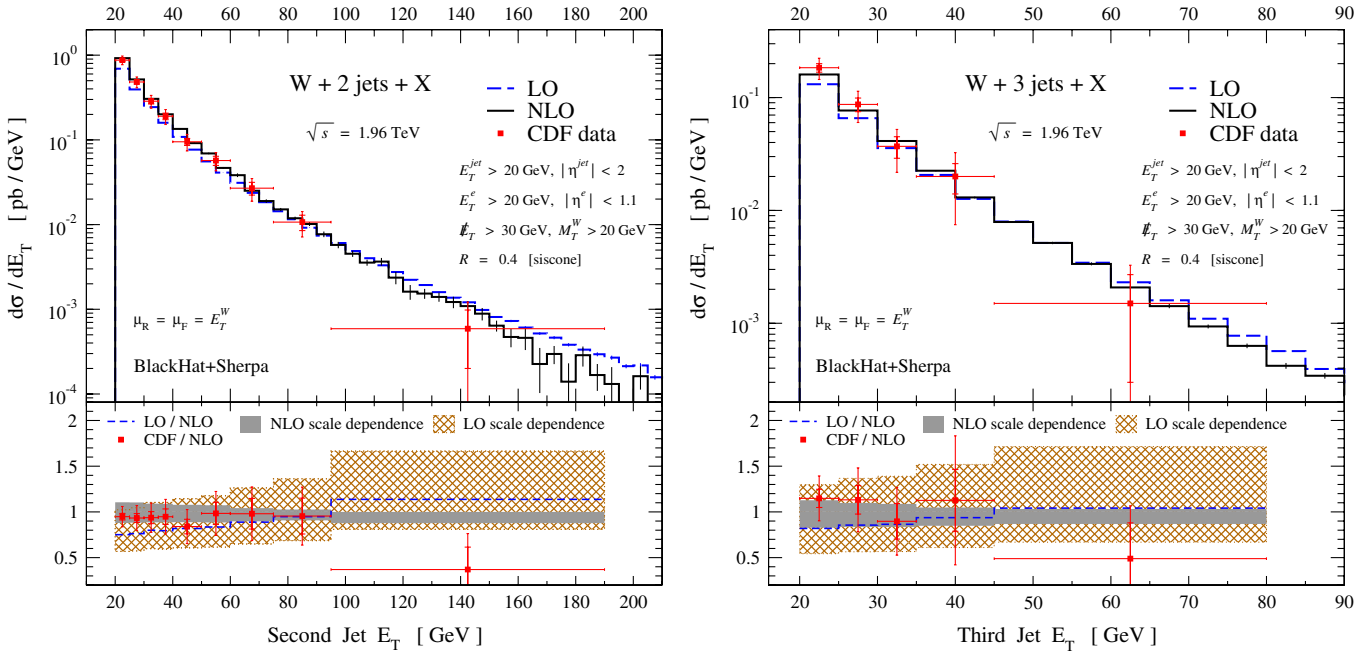

FIG. 5 (color online). The measured cross section, $d\sigma(p\bar{p} \to e\nu + \geq n\text{-jets})/dE_T^{n\text{th-jet}}$, for inclusive $W + n$-jet production, compared to full NLO predictions for $n = 2, 3$. In the upper panels the NLO distribution is the solid (black) histogram, and CDF data points are the (red) points, whose inner and outer error bars, respectively, denote the statistical and total uncertainties (excluding the luminosity error) on the measurements added in quadrature. The LO predictions are shown as dashed (blue) lines. The thin vertical lines in the center of each bin (where visible) give the numerical integration errors for that bin. Each lower panel shows the distribution normalized to the full NLO prediction, using the CDF experimental bins (that is, averaging over bins in the upper panel). The scale-dependence bands are shaded (gray) for NLO and crosshatched (brown) for LO.

dences still remain. (These calculations make different choices for the scale variation and are not directly comparable to the LO parton-level predictions given here.) As in the experimental analysis, we sum the W^- and W^+ cross sections, which are identical at the Tevatron (for forward-backward symmetric acceptance cuts).

We have also computed the $W + 2$-jet and $W + 3$-jet total cross sections at NLO with a larger merging parameter, $f = 0.75$. (CDF uses a value of $f = 0.75$ [35], but for a different, infrared-unsafe algorithm, JETCLU.) The value of the NLO $W + 3$-jet production cross section of 0.882 pb in Table II then increases to 0.917 pb (about 4%). The $W +$ 2-jet production cross section shows a more modest increase from 7.62 pb to 7.69 pb (about 1%). Distributions, as, for example, the ones shown in Fig. 5 (see also Table III), follow a similar bin-by-bin dependence on f.

In Fig. 5, we compare the E_T distribution of the second- and third-most energetic jets in CDF data [35] to the NLO predictions for $W + 2$-jet and $W + 3$-jet production, respectively. For convenience, in Table III we collect the data used to construct the third-jet E_T plot in Fig. 5. We include scale-dependence bands obtained as described above.[5] The experimental statistical and systematic uncertainties (excluding an overall luminosity uncertainty of 5.8%) have been combined in quadrature. The upper panels of Fig. 5 show the distribution itself, while the lower panels show the ratio of the LO value and of the data to the NLO result for the central value of $\mu = E_T^W$. Note that we normalize here to the NLO result, not to LO as done elsewhere. The LO/NLO curve in the bottom panel represents the inverse of the so-called K factor (NLO to LO ratio).

We do not include PDF uncertainties in our analysis. For $W + 1, 2$-jet production at the Tevatron these uncertainties have been estimated in Ref. [35]. For these processes, they

[5]We emphasize that the scale-uncertainty bands are only rough estimates of the theoretical error, which would properly be given by the difference between an NLO result and one to higher order (next-to-next-to-leading order).

TABLE III. The differential cross sections, $d\sigma(p\bar{p} \to e\nu + \geq 3\text{-jets})/dE_T^{3\text{rd-jet}}$, for $W + 3$-jet production at the Tevatron using the experimental cuts (3.1) and (3.2) of Ref. [35]. This table corresponds to the values plotted in Fig. 5.

$E_T^{3\text{rd-jet}}$	$d\sigma/dE_T^{3\text{rd-jet}}$ (pb/GeV)		
	CDF	LO	NLO
20–25	0.184 ± 0.0394	$0.131^{+0.0769}_{-0.0443}$	$0.160^{+0.0205}_{-0.0277}$
25–30	0.087 ± 0.0268	$0.066^{+0.0393}_{-0.0224}$	$0.077^{+0.0075}_{-0.0126}$
30–35	0.037 ± 0.0153	$0.036^{+0.0216}_{-0.0123}$	$0.041^{+0.0036}_{-0.0068}$
35–45	0.020 ± 0.0125	$0.017^{+0.0103}_{-0.0058}$	$0.018^{+0.0009}_{-0.0027}$
45–80	$0.0015 \pm 0.001\,77$	$0.0032^{+0.002\,07}_{-0.001\,14}$	$0.0031^{+0.000\,15}_{-0.000\,41}$

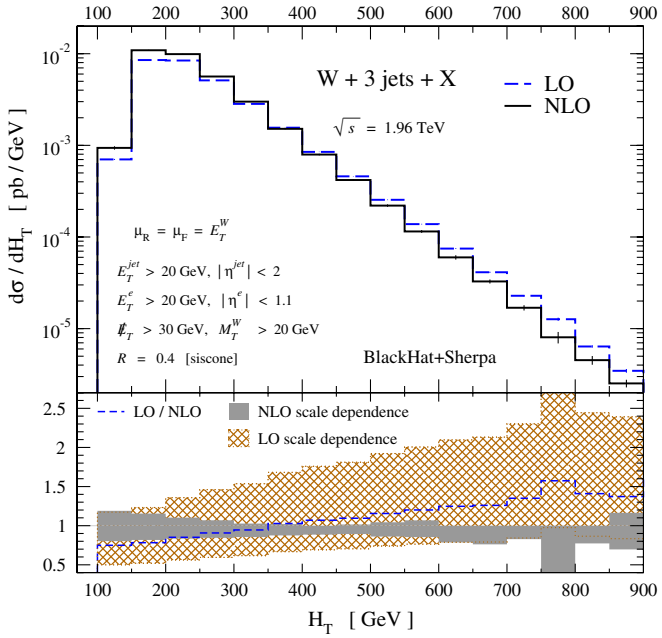

FIG. 6 (color online). Theoretical predictions for the H_T distribution in $W + 3$-jet production at the Tevatron. The LO prediction is shown by a dashed (blue) line; the NLO one by a solid (black) one. The upper panel shows the distribution itself, while the lower panel shows the distributions and scale-dependence bands, crosshatched (brown) for LO and solid (gray) for NLO, normalized to the NLO prediction for $\mu = E_T^W$. The numerical integration errors, indicated by thin vertical lines in the upper panel, are noticeable only in the tail.

are smaller than uncertainties associated with NLO scale dependence at low jet E_T, but larger at high E_T.

For reference, we also show the LO distributions and corresponding scale-dependence bands. The NLO predictions match the data very well, and uniformly (without any difference in slope) in all but the highest E_T experimental bin. The central values of the LO predictions, in contrast, have different shapes from the data. In the upper panels, we have used 5 GeV bins to plot the predictions, and have superposed the data points, although CDF used different bins in their analysis. In the lower panel, which shows the ratio of the LO prediction, and of the data, to the NLO prediction, we have used the experimental bins, which are wider at higher E_T. A very similar plot was given previously [34], based on a particular leading-color approximation. As we discuss in Sec. VIII, those results differ only slightly from the complete NLO results presented here.

In Fig. 6, we show the distribution for the total transverse energy H_T, given in Eq. (2.14). This quantity has been used in top-quark studies, and will play an important role in searches for decays of heavy new particles at the LHC. The upper panel shows the LO and NLO predictions for the distribution, and the lower panel their ratio. The NLO scale-dependence band, as estimated using five points, ranges from $\pm 20\%$ around its central value at low H_T to $\pm 5\%$ around 400 GeV, and back to around $\pm 10\%$ at 800 GeV. The band is accidentally narrow at energies near the middle of graph, because the curves associated with the five μ values converge as the H_T value rises from lower values towards the middle ones. (The fluctuations visible in the tail of the distribution are a reflection of the limited statistics for the Monte Carlo integration, as we show a larger dynamical range than in the E_T spectrum.) The shape of the LO distribution is noticeably different, for any of the μ values, from that at NLO. At low H_T, the central LO prediction is 20% below the NLO central value, whereas at the largest H_T it is nearly 50% higher. Thus for $\mu = E_T^W$ the NLO correction cannot be characterized by a

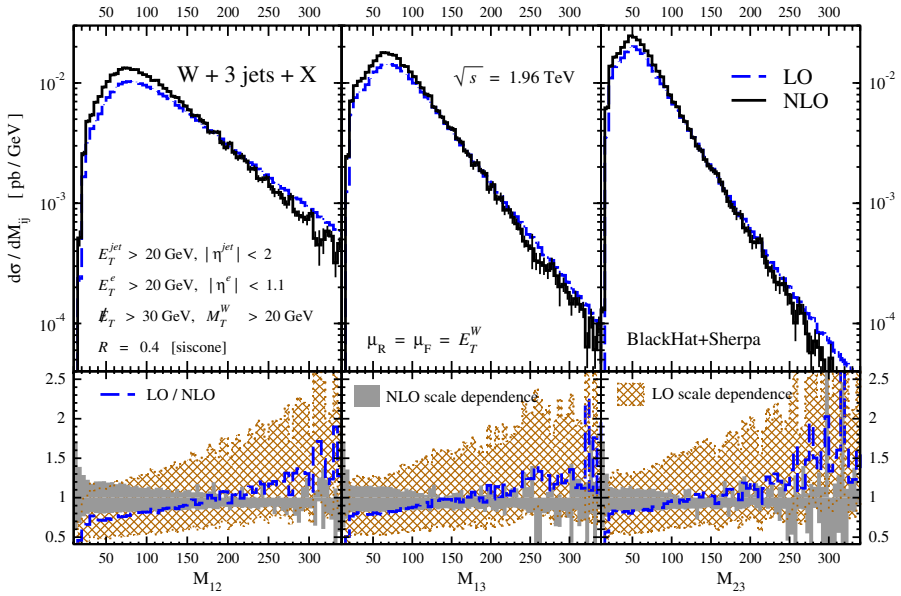

FIG. 7 (color online). LO and NLO predictions for the dijet invariant mass distributions M_{ij} (in GeV) in $W + 3$-jet production at the Tevatron. The histograms and bands have the same meaning as in Figs. 5 and 6.

constant K factor (ratio of NLO to LO results). We will address some of the reasons for the difference in shape in the following section. We note that the NLO scale-dependence band has a somewhat different appearance from the corresponding figure in Ref. [34], because the latter used wider bins at large H_T and had larger integration errors.

In Fig. 7, we show the distributions for the three dijet invariant masses we can form: hardest and middle jet M_{12}, hardest and softest jet M_{13}, and middle and softest jet M_{23}. The NLO scale-dependence bands are somewhat broader than for the E_T or H_T distributions. The distributions become increasingly steep as we move from masses of hardest to softer jets. That gross feature is unaltered in passing from LO to NLO, although each distribution falls off somewhat faster at NLO, as was the case for H_T.

IV. CHOOSING SCALES

The renormalization and factorization scales are not physical scales. As such, physical quantities should be independent of them. They arise in theoretical calculations as artifacts of defining α_S and the parton distributions, respectively. We will follow the usual practice and choose the two to be equal, $\mu_R = \mu_F = \mu$. The sensitivity of a perturbative result to the common scale is due to the truncation of the perturbative expansion; this dependence would be canceled by terms at higher orders. NLO calculations greatly reduce this dependence compared to LO results, but of course do not eliminate it completely. In practice, we must therefore choose this scale. Intuitively, we would expect a good choice for μ to be near a ''characteristic'' momentum scale p for the observable we are computing, in order to minimize logarithms in higher-order terms of the form $\ln(\mu/p)$. The problem is that complicated processes such as $W + 2$, 3-jet production have many intrinsic scales, and it is not clear we can distill them into a single number. For any given point in the fully-differential cross section, there is a range of scales one could plausibly choose. One could choose a fixed scale μ, the same for all events. However, because there can be a large dynamic range in momentum scales (particularly at the LHC, where jet transverse energies well above M_W are not uncommon), it is natural to pick the scale μ dynamically, on an event-by-event basis, as a function of the observable or unobservable parameters of an event.

A particularly good choice of scale might minimize changes in shape of distributions from LO to NLO, such as those visible in Figs. 6 and 7. Such a choice might in turn make it possible for LO programs incorporating parton showering and hadronization [7,8,84] to be more easily reweighted to reflect NLO results.

Before turning to dynamical scales and kinematic distributions, let us first examine how the total cross section depends on a fixed scale. In Fig. 8 we display this dependence for the Tevatron[6] and the LHC (left and right, respectively). We vary the scale between $\mu = M_W/2$ and $4M_W$ for the Tevatron, and between $M_W/2$ and $16M_W$ for the LHC. We must be careful to vary the scale in a ''sensible'' range. For the NLO calculation, in particular, we do not wish to reintroduce large logarithms of scales. The figure shows the characteristic increasing-and-decreasing of the NLO prediction (see e.g. Refs. [85,86]) as well as the monotonicity of the LO one. It also shows a substantial reduction in scale dependence going from LO to NLO. The lower panels show the K factor. The large sensitivity of the LO cross section to the choice of scale implies a similar large dependence in this ratio.

We thus see that, as expected, the total cross sections at NLO are much less sensitive to variations of the scale than at LO. We now turn to the scale dependence of kinematic distributions. In this case the K factor will not only be sensitive to the scale chosen, but it will in general depend on the kinematic variable. We will see that a poor choice of scale can lead to problems not only at LO, but also at NLO, especially in the tails of distributions.

The sensitivity to a poor scale choice is already noticeable at the Tevatron, in the shape differences between LO and NLO predictions visible in Figs. 6 and 7. However, it becomes more pronounced at the LHC because of the larger dynamical range of available jet transverse energies. We can diagnose particularly pathological choices of scale using the positivity of the NLO cross section: too low a scale at NLO will make the total cross section unphysically negative.

This diagnostic can be applied bin by bin in distributions. For example, in Fig. 9 we show the E_T distribution of the second-most energetic jet of the three, at the LHC. In the left plot we choose the scale to be the W transverse energy E_T^W [defined in Eq. (3.3)] used earlier in the Tevatron analysis. Near an E_T of 475 GeV, the NLO prediction for the differential cross section turns negative! This is a sign of a poor scale choice, which has reintroduced large enough logarithms of scale ratios to overwhelm the LO terms at that jet E_T. Its inadequacy is also indicated by the large ratio of the LO to NLO distributions at lower E_T, and in the rapid growth of the NLO scale-dependence band with E_T. In contrast, the right panel of Fig. 9 shows that $\hat{H}_T$ [defined in Eq. (2.10)] provides a sensible choice of scale: the NLO cross section stays positive, and the ratio of the LO and NLO distributions, though not completely flat, is much more stable.

Why is $\mu = E_T^W$ such a poor choice of scale for the second jet E_T distribution, compared with $\mu = \hat{H}_T$? (For an independent, but related discussion of this question, see Ref. [40].) Consider the two distinct types of $W + 3$ jet configurations shown in Fig. 10. If configuration (a) domi-

[6]Note that the Tevatron plot is for $E_T^{\text{3rd-jet}} > 20$ GeV, not the cut $E_T^{\text{3rd-jet}} > 25$ GeV used in Table II.

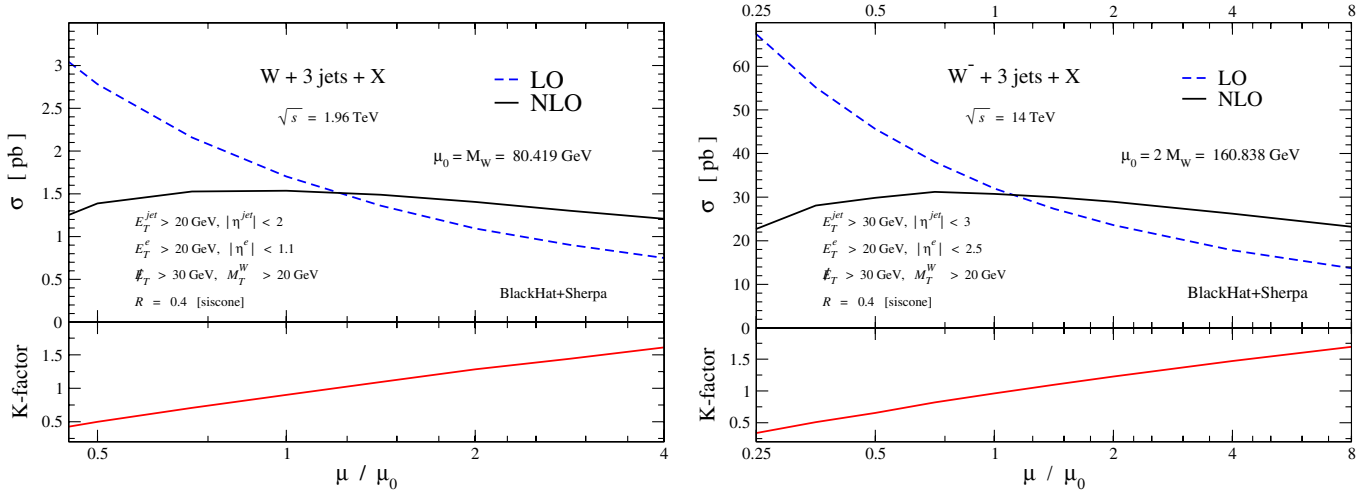

FIG. 8 (color online). The scale dependence of the cross section for W + 3-jet production at the Tevatron is shown on the left, as a function of the common renormalization and factorization scale μ, with $\mu_0 = M_W$. Similar curves for W^- + 3-jet production at the LHC are shown on the right, except that $\mu_0 = 2M_W$. In each case, the lower panel shows the K factor.

nated, then as the jet E_T increased, E_T^W would increase along with it, by conservation of transverse momentum. However, in configuration (b), the bulk of the transverse momentum can be balanced between the first and second jet, with the W and the third jet remaining soft. In the tail of the second-jet E_T distribution, configuration (b) is highly favored kinematically, because it implies a much smaller partonic center-of-mass energy. Because E_T^W remains small, the wrong scale is being chosen in the tail. Evidence for the dominance of configuration (b) over (a) in W + 2-jet production can be found in Ref. [40], which shows that the two jets become almost back to back as the jet E_T cut rises past M_W. The negative NLO cross section in the left panel of Fig. 9 provides evidence of the same domination in W + 3-jet production.

However, configuration (b) also tends to dominate in the tails of generic multijet distributions, such as H_T or M_{ij}, in which large jet transverse energies are favored. The reason is that for jet transverse energies well above M_W, the W behaves like a massless vector boson, and so there is a

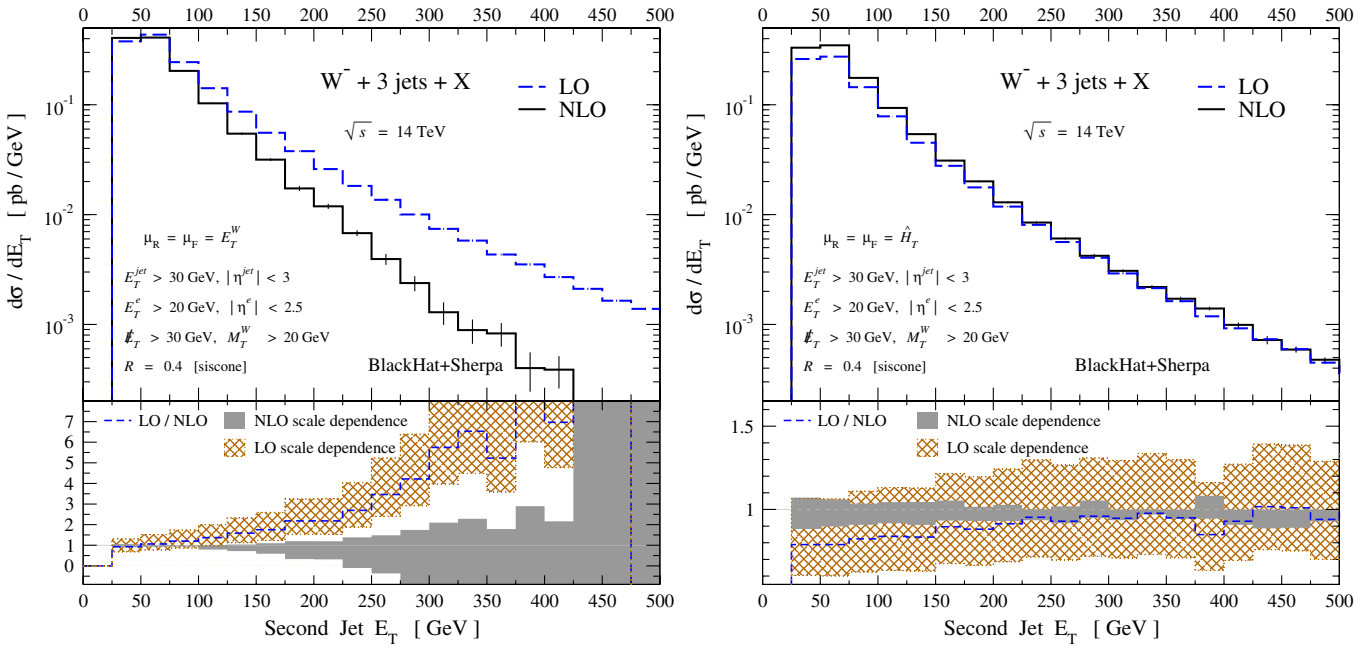

FIG. 9 (color online). The E_T distribution of the second jet at LO and NLO, for two dynamical scale choices, $\mu = E_T^W$ (left plot) and $\mu = \hat{H}_T$ (right plot). The histograms and bands have the same meaning as in previous figures. The NLO distribution for $\mu = E_T^W$ turns negative beyond $E_T = 475$ GeV.

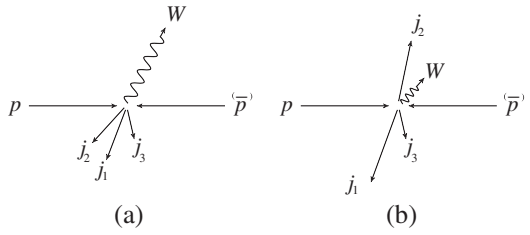

FIG. 10. Two distinct $W + 3$ jet configurations with rather different scales for the W transverse energy. In configuration (a) an energetic W balances the energy of the jets, while in (b) the W is relatively soft. Configuration (b) generally dominates over (a) when the jet transverse energies get large.

kinematic enhancement when it is soft, as in configuration (b). Exceptions would be in distributions such as the transverse energy of the W itself, or of its decay lepton, which kinematically favor configuration (a) in their tails.

In contrast to $\mu = E_T^W$, the scale $\mu = \hat{H}_T$ becomes large in the tails of generic multijet transverse-energy distributions. For the distribution of the second jet E_T, this is evident from the close agreement between LO and NLO values, shown in the right panel of Fig. 9. The same features are evident, though less pronounced, in the H_T distributions shown in Fig. 11. The left plot is again for $\mu = E_T^W$, and the right plot for $\mu = \hat{H}_T$. The shapes of the LO and NLO distributions for $\mu = E_T^W$ are quite different; the ratio displayed varies from around 1 at H_T of 200 GeV to around 2 at H_T near 1200 GeV. In contrast, the ratio for $\mu = \hat{H}_T$ is nearly flat.

These features are not special to the H_T distribution itself. For example, Fig. 12 displays the ratio of LO to NLO predictions for two other $W + 3$-jet distributions for the two scale choices. The left panel shows the ratios for the leading dijet mass, while the right panel shows ratios for the leading ΔR distribution. Once again the ratios for $\mu = \hat{H}_T$ have a much milder dependence than those for $\mu = E_T^W$.

As we shall see further in the next section, the roughly flat ratio for the choice $\mu = \hat{H}_T$ holds for a wide variety of distributions. It does not hold for all: some NLO corrections cannot be absorbed into a simple redefinition of the renormalization scale. The distribution of the second-most energetic jet in Fig. 9 provides one example. A second example, discussed below, is the H_T distribution for $W^+ +$ 2-jet production in the left plot of Fig. 15. A third example (not shown) would be the H_T distribution for $W + 1$-jet production; this case is easy to understand because only configuration (a) (with the second and third jets erased) is available at LO, while configuration (b) can dominate at NLO, so effectively a new subprocess opens up at NLO.

Although the E_T^W scale choice is a poor one as far as the tails of many distributions are concerned, we note that it does give reasonable results for the Tevatron and LHC total cross sections with our standard jet cuts, which are dominated by modest jet transverse energies. For $\mu = E_T^W$, the NLO cross section for $W^- + 3$-jet production at the LHC is $31.37(0.20)^{+0.0}_{-2.47}$ pb, which has much smaller scale variation than the LO result $37.16(0.07)^{+16.35}_{-10.35}$ pb. (The parentheses indicate the integration uncertainties, and subscripts

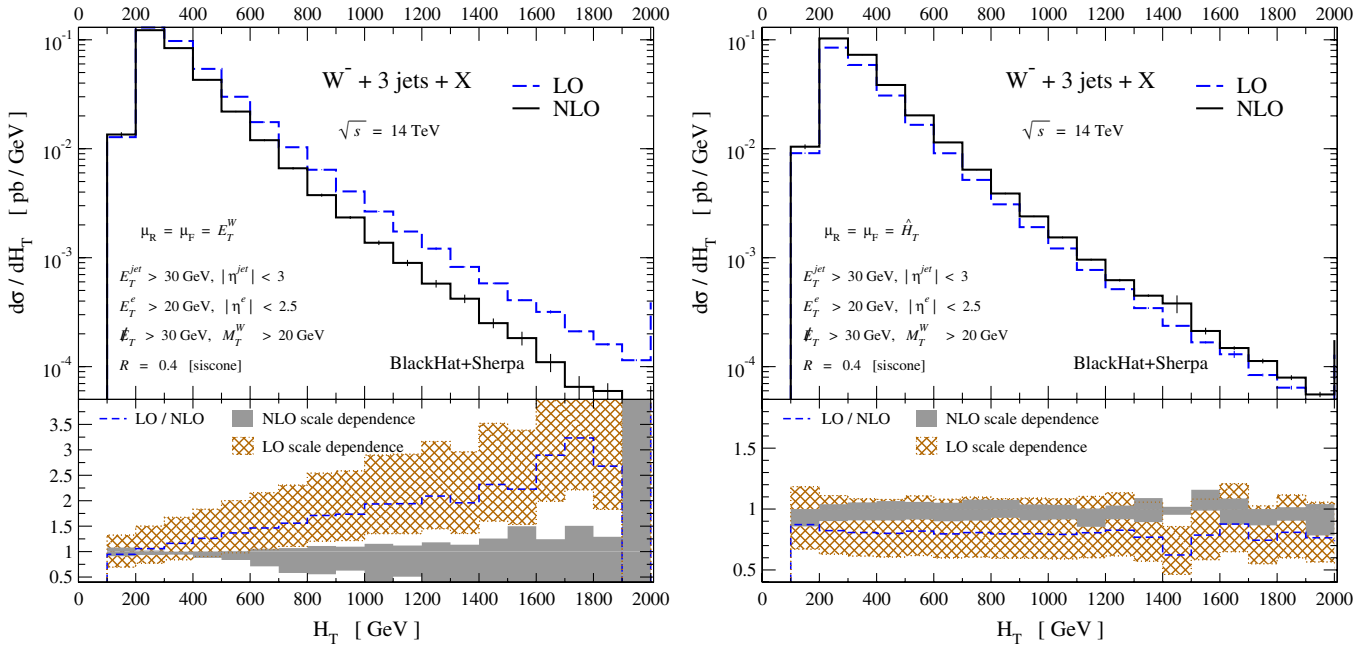

FIG. 11 (color online). The H_T distribution for $W^- + 3$-jet production at the LHC. The scale choices $\mu = E_T^W$ and $\mu = \hat{H}_T$ are shown, respectively, on the left and the right. The histograms and bands have the same meaning as in previous figures.

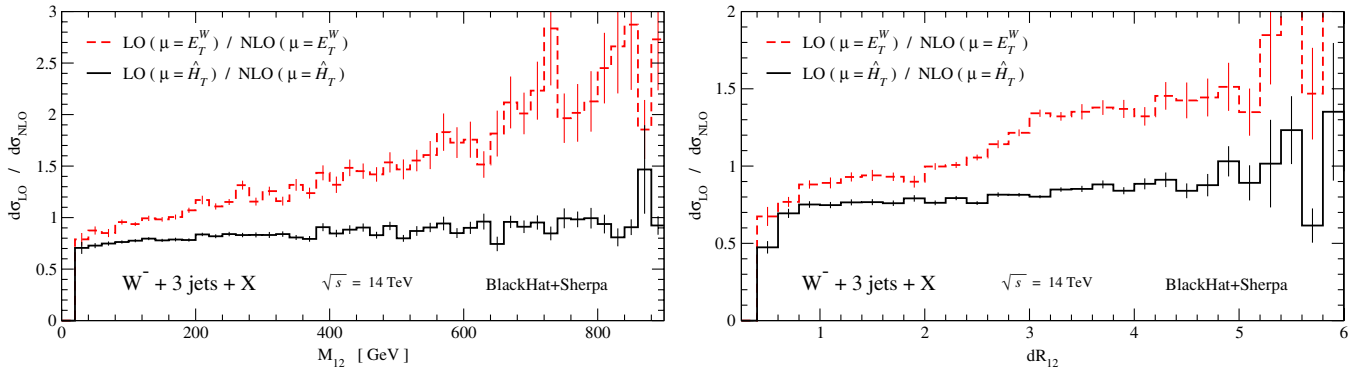

FIG. 12 (color online). Ratios of LO to NLO predictions for the distributions in the dijet invariant mass (left panel) and ΔR separation (right panel) for the leading two jets in W^- + 3-jet production at the LHC. In each panel, the dashed (red) line gives the scale choice $\mu = E_T^W$, while the solid (black) line gives the (much flatter) ratio for $\mu = \hat{H}_T$.

and superscripts the scale variation.) For $\mu = \hat{H}_T$, the NLO value is $27.52(0.14)^{+1.34}_{-2.81}$ pb; the two NLO results are consistent within the scale variation band.

Accordingly, to have a proper description of distributions, we adopt $\hat{H}_T$ as our default choice of scale for W + 3-jet production at the LHC. In Fig. 13 we display the scale variation of the total cross section, evaluating it at the five scales $\mu_0/2$, $\mu_0/\sqrt{2}$, μ_0, $\sqrt{2}\mu_0$, $2\mu_0$ with $\mu_0 = \hat{H}_T$. As usual, the variation is much smaller at NLO than at LO. Because $\hat{H}_T$ includes a scalar sum, it is somewhat larger than an "average" momentum transfer. One could choose a scale lower by a fixed ratio, say $\hat{H}_T/2$. This would shift the LO-to-NLO ratio curves in Figs. 11 and 12, for example, up towards a ratio of 1. It would have only a modest effect on the NLO predictions, however, because the scale-dependence curve for the NLO cross section is relatively flat.

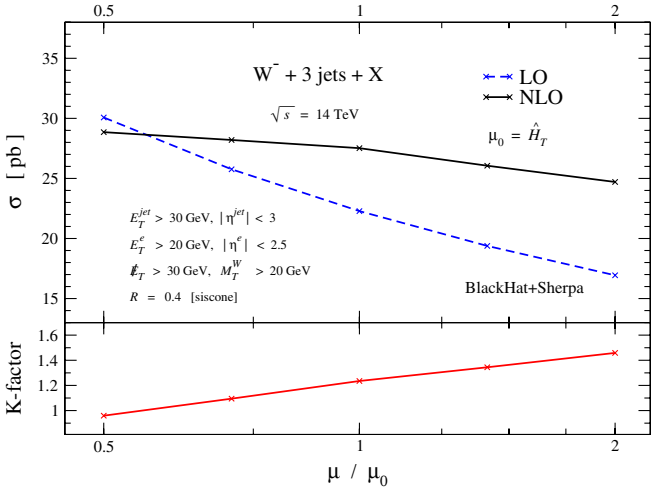

FIG. 13 (color online). The scale dependence of the cross section for W^- + 3-jet production at the LHC as a function of μ/μ_0, with $\mu_0 = \hat{H}_T$.

It is interesting to compare our default choice $\mu = \hat{H}_T$ with the choice of scale advocated in Ref. [40] on the basis of soft-collinear effective theory,

$$\mu_{\rm had}^2 = \tfrac{1}{4}M_{\rm had}^2 + M_W^2. \tag{4.1}$$

In this equation $M_{\rm had}$ is the invariant mass of the jets. (As explained in Ref. [40], the factor of 1/4 is a choice, not dictated by a principle.) With this choice, one can greatly reduce the shift between LO and NLO in W + 2-jet distributions, compared to more conventional choices such as $\mu = E_T^W$. We have confirmed that for W + 2-jet production with a few exceptions, such as the decay lepton transverse energy, the choice $\mu = \mu_{\rm had}$ does fare better than $\mu = E_T^W$ in bringing LO in line with NLO. How does this choice fare in W + 3-jet production? To answer this question, we have compared several distributions. In Fig. 14, we consider the E_T distributions of the first, second and third jets in W^+ + 3-jet production at the LHC. We compare the LO results for $\mu = \hat{H}_T$ and $\mu = \mu_{\rm had}$ to the reference NLO results for $\mu = \hat{H}_T$. (Any sensible choice of scale at NLO should give very similar results.) As can be seen from the figure, the choice $\mu = \hat{H}_T$ leads to a somewhat flatter LO to NLO ratio than does $\mu = \mu_{\rm had}$ for the first jet, and performs about as well for the second and third leading jets.

It is also instructive to compare the H_T distributions for W + 2-jet and W + 3-jet production. The left panel of Fig. 15 shows the distribution in W^+ + 2-jet production; here the scale $\mu = \mu_{\rm had}$ gives an LO result closer to the NLO one. On the other hand, in the right panel, which shows the distribution in W^+ + 3-jet production, the choice $\mu = \hat{H}_T$ gives a LO to NLO ratio which is comparably flat to the $\mu = \mu_{\rm had}$ choice.

In contrast, examine the positron p_T (or E_T) distribution, shown in Fig. 16 for W^+ + 2, 3-jet production at the LHC. As can be seen in the lower panel of each plot, the choice $\mu = \hat{H}_T$ performs better than $\mu = \mu_{\rm had}$ at LO, in match-

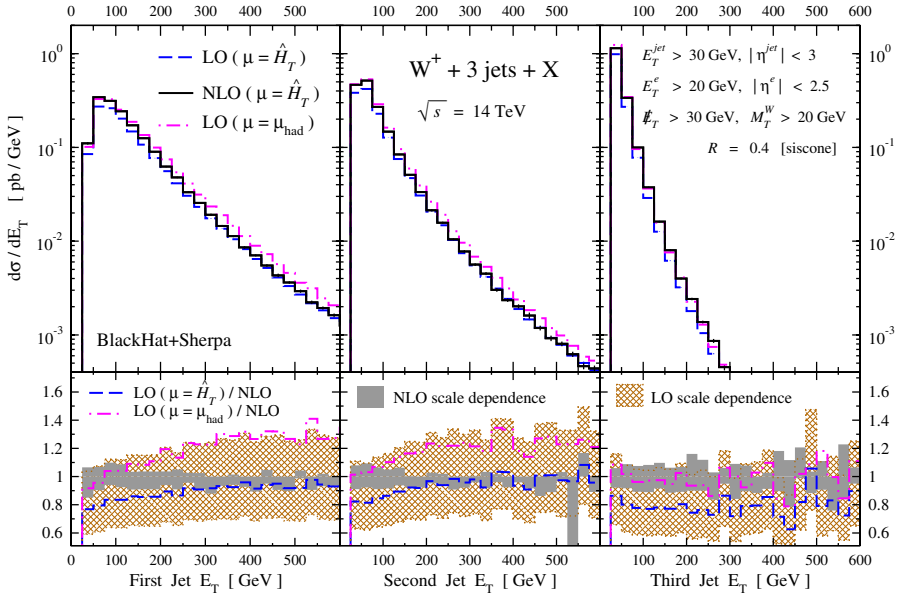

FIG. 14 (color online). The first, second, and third jet E_T distributions for $W^+ + 3$-jet production. The scale choice $\mu = \mu_{\rm had}$ at LO, shown by the dot-dashed (magenta) lines, is compared with the LO and NLO results using our default scale choice $\mu = \hat{H}_T$, shown, respectively, in the dashed (blue) and solid (black) lines. The scale-dependence bands are shaded (gray) for NLO and crosshatched (brown) for LO.

ing the more accurate NLO result at large values of E_T^e. The reason is that large E_T^e forces the W transverse energy to be large, which in turn favors configuration (a) in Fig. 10, in which a relatively low-mass cluster of jets recoils against the W boson. Thus the scale $\mu = \mu_{\rm had}$ drops below the typical momentum transfer in the process.

In summary, both $\mu = \mu_{\rm had}$ and $\mu = \hat{H}_T$ are a great improvement over the scale choice $\mu = E_T^W$. For some distributions $\mu = \mu_{\rm had}$ is a somewhat better choice at LO than $\mu = \hat{H}_T$ while for other distributions $\mu = \hat{H}_T$ is better. These attributes should not come as a surprise, given the multiscale nature of jet production.

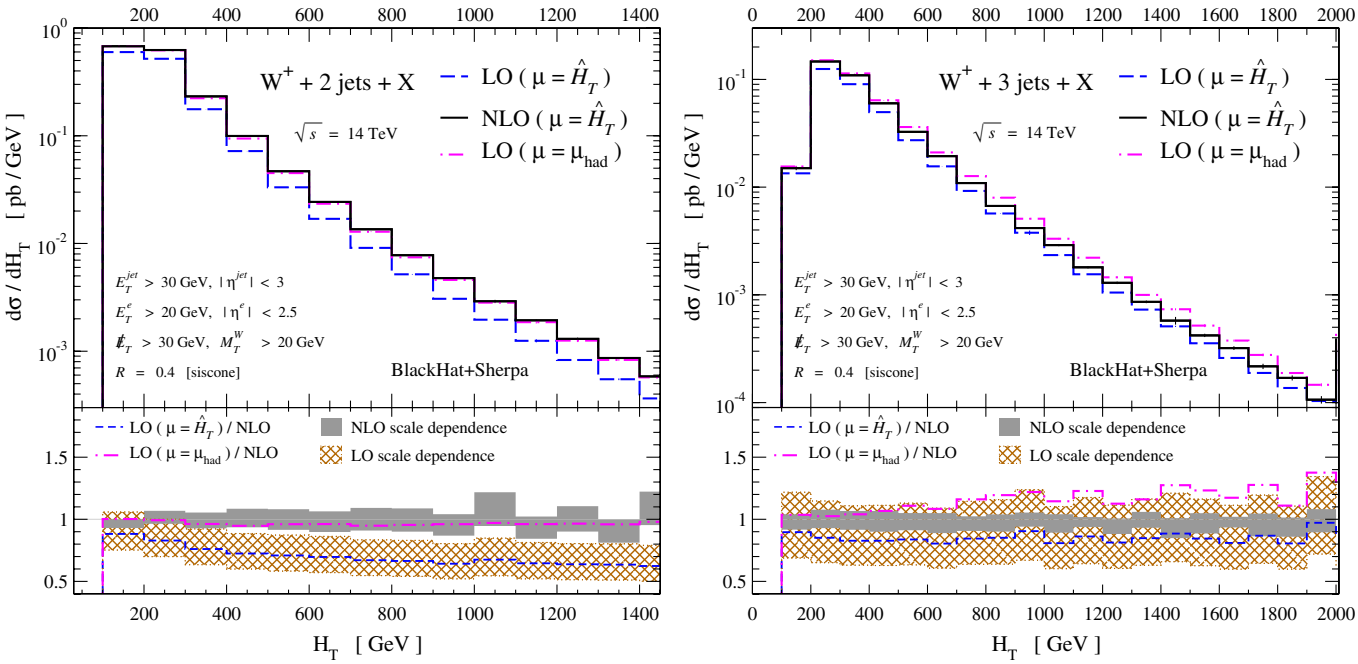

FIG. 15 (color online). The H_T distributions for $W^+ + 2$, 3-jet production. LO results for $\mu = \mu_{\rm had}$ are compared with LO and NLO results for $\mu = \hat{H}_T$. The lines and bands have the same meaning as in Fig. 14.

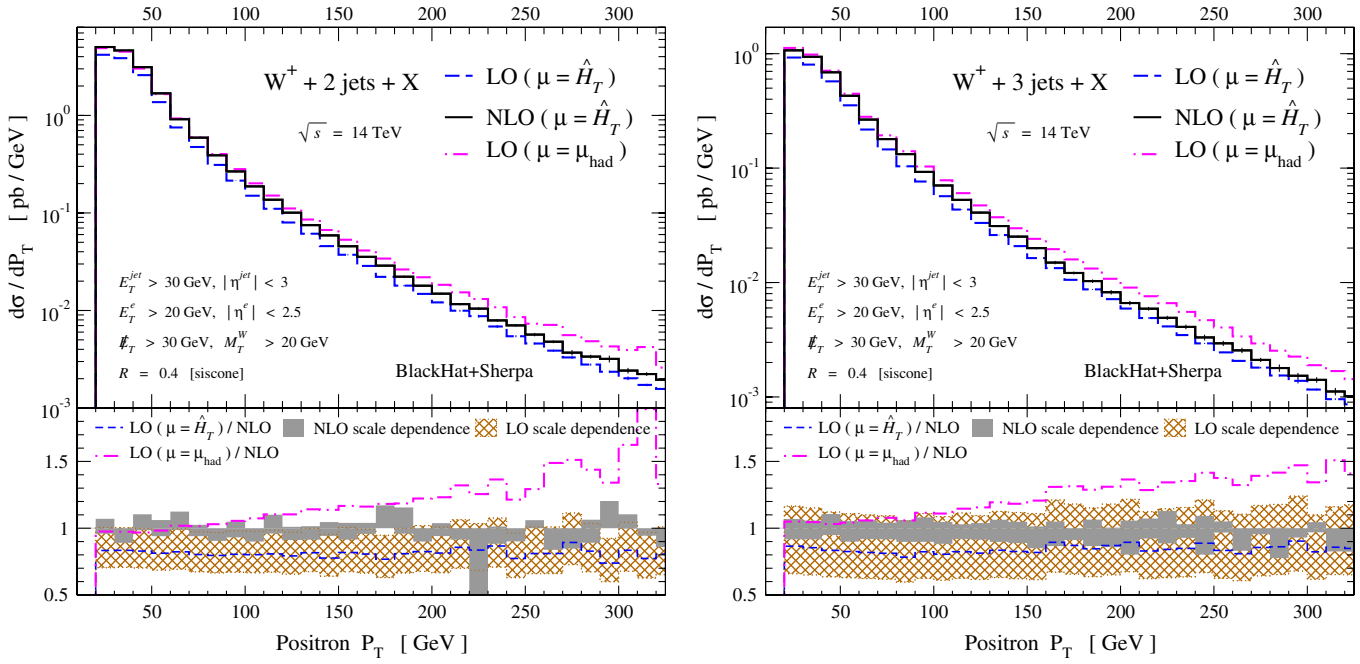

FIG. 16 (color online). The distribution in the positron transverse momentum for $W^+ + 2$, 3-jet production at the LHC. LO results for $\mu = \mu_{\rm had}$ are compared with LO and NLO results for $\mu = \hat{H}_T$. The lines and bands have the same meaning as in Fig. 14.

V. PREDICTIONS FOR THE LHC

In this section we present the first complete NLO predictions for $W + 3$-jet production at the LHC. The initial run of the LHC will almost certainly not be at its full design energy of 14 TeV, but we choose this energy to simplify comparisons to earlier studies. Most of the features visible at 14 TeV would of course remain at the lower energy, such as 10 TeV, of an initial run. The production of $W + 3$ jets at the LHC was also studied at NLO in Ref. [41], however with a set of subprocesses accounting for only 70% of the cross section; for on-shell W bosons; and with a less accurate leading-color approximation than that of Ref. [34]. For our analysis of $W + 3$-jet production at the LHC, we use the following kinematical cuts,

$$|\eta^{\rm jet}| < 3, \qquad R = 0.4, \qquad |\eta^e| < 2.5,$$
$$E_T^e > 20\ {\rm GeV}, \qquad E_T^\nu > 30\ {\rm GeV}, \qquad M_T^W > 20\ {\rm GeV}. \tag{5.1}$$

We also quote total cross sections with both of the following jet cuts

$$E_T^{\rm jet} > 30\ {\rm GeV} \quad \text{and} \quad E_T^{\rm jet} > 40\ {\rm GeV}. \tag{5.2}$$

We show distributions only using the first of these two cuts. We employ the SISCONE jet algorithm [37] everywhere (with f parameter set to 0.5), except for Tables VI and VII where we use the k_T algorithm [87].

For the LHC we adopt the default factorization and renormalization scale choices,

$$\mu = \hat{H}_T, \tag{5.3}$$

where $\hat{H}_T$ is defined in Eq. (2.10). As discussed in the previous section, this choice does not have the shortcomings of $\mu = E_T^W$ in describing the large transverse energy tails of generic distributions.

At the LHC, a pp collider, the total rates and the shapes of some distributions are quite different for W^- and W^+ production. At 14 TeV, the qg initial state accounts for over half of $W + n$-jet production. There are considerably more u quarks than d quarks in the proton in the relevant range of the momentum fraction x, leading to greater production of W^+ than W^-. Accordingly, we quote separate results for total cross sections in Tables IV, V, VI, and VII. In Table IV, we show the $W^- + 1$, 2, 3-jet cross sections using the SISCONE algorithm, for two different choices of jet E_T cut, 30 and 40 GeV. The corresponding results for $W^+ + 1$, 2, 3-jet production are given in Table V. In Tables VI and VII, we show the corresponding results for the k_T jet algorithm with a pseudocone radius of 0.4, for W^- and W^+ production, respectively. It is interesting to note that while the NLO cross sections for $W + 1$, 2-jet production are larger for the SISCONE algorithm than for k_T (with the algorithm parameters we have chosen), the relative size is reversed for $W + 3$-jet production. (The entries for the LO $W + 1$-jet cross section are identical for the SISCONE and k_T algorithms because the same set of events was used to compute them.)

We next describe NLO results for kinematic distributions. For distributions that do not differ appreciably for

TABLE IV. Cross sections for W^- production using the SISCONE jet algorithm, with jet cuts $E_T^{\rm jet} > 30$ GeV or $E_T^{\rm jet} > 40$ GeV. The remaining cuts are as in Eq. (5.1).

Number of jets	LO $E_T^{\rm jet} > 30$ GeV	NLO $E_T^{\rm jet} > 30$ GeV	LO $E_T^{\rm jet} > 40$ GeV	NLO $E_T^{\rm jet} > 40$ GeV
1	$343.29(0.18)^{+15.65}_{-15.43}$	$456.60(1.43)^{+16.61}_{-10.10}$	$215.68(0.12)^{+12.19}_{-11.33}$	$298.44(0.77)^{+12.75}_{-8.43}$
2	$99.78(0.09)^{+20.81}_{-15.60}$	$122.71(0.92)^{+5.88}_{-7.41}$	$58.52(0.063)^{+12.49}_{-9.41}$	$72.96(0.54)^{+3.20}_{-4.54}$
3	$22.28(0.04)^{+7.80}_{-5.34}$	$27.52(0.14)^{+1.34}_{-2.81}$	$11.012(0.02)^{+3.87}_{-2.67}$	$13.96(0.07)^{+1.03}_{-1.31}$

TABLE V. Cross sections for W^+ production using SISCONE, with jet cut $E_T^{\rm jet} > 30$ GeV or $E_T^{\rm jet} > 40$ GeV. The remaining cuts are as in Eq. (5.1).

Number of jets	LO $E_T^{\rm jet} > 30$ GeV	NLO $E_T^{\rm jet} > 30$ GeV	LO $E_T^{\rm jet} > 40$ GeV	NLO $E_T^{\rm jet} > 40$ GeV
1	$469.37(0.32)^{+21.86}_{-21.26}$	$615.77(2.04)^{+23.76}_{-14.39}$	$301.20(0.22)^{+17.06}_{-15.86}$	$415.50(1.90)^{+19.40}_{-12.86}$
2	$143.91(0.18)^{+29.92}_{-22.43}$	$174.28(0.48)^{+6.56}_{-10.37}$	$86.32(0.12)^{+18.33}_{-13.81}$	$105.99(0.31)^{+5.36}_{-5.82}$
3	$34.75(0.05)^{+12.06}_{-8.31}$	$41.47(0.27)^{+2.81}_{-3.50}$	$17.64(0.02)^{+6.14}_{-4.25}$	$21.76(0.15)^{+1.68}_{-1.86}$

TABLE VI. Cross sections for W^- production using the k_T algorithm ($R = 0.4$), with jet cut $E_T^{\rm jet} > 30$ GeV or $E_T^{\rm jet} > 40$ GeV. The remaining cuts are as in Eq. (5.1).

Number of jets	LO $E_T^{\rm jet} > 30$ GeV	NLO $E_T^{\rm jet} > 30$ GeV	LO $E_T^{\rm jet} > 40$ GeV	NLO $E_T^{\rm jet} > 40$ GeV
1	$343.29(0.18)^{+15.65}_{-15.43}$	$444.75(1.44)^{+15.12}_{-8.85}$	$215.68(0.12)^{+12.19}_{-11.33}$	$290.44(0.77)^{+11.65}_{-7.55}$
2	$102.88(0.09)^{+21.40}_{-16.05}$	$120.07(0.86)^{+4.19}_{-6.33}$	$59.99(0.06)^{+12.78}_{-9.63}$	$70.85(0.42)^{+2.12}_{-3.87}$
3	$25.84(0.05)^{+8.99}_{-6.17}$	$29.29(0.16)^{+0.65}_{-2.32}$	$12.78(0.02)^{+4.46}_{-3.09}$	$14.89(0.08)^{+0.59}_{-1.18}$

TABLE VII. Cross sections for W^+ production using the k_T algorithm, with jet cut $E_T^{\rm jet} > 30$ GeV or $E_T^{\rm jet} > 40$ GeV. The remaining cuts are as in Eq. (5.1).

Number of jets	LO $E_T^{\rm jet} > 30$ GeV	NLO $E_T^{\rm jet} > 30$ GeV	LO $E_T^{\rm jet} > 40$ GeV	NLO $E_T^{\rm jet} > 40$ GeV
1	$469.37(0.32)^{+21.86}_{-21.26}$	$600.66(2.06)^{+21.83}_{-12.82}$	$301.20(0.22)^{+17.06}_{-15.83}$	$405.27(1.91)^{+17.91}_{-11.82}$
2	$148.46(0.19)^{+30.78}_{-23.08}$	$171.45(0.50)^{+3.81}_{-9.39}$	$88.48(0.12)^{+18.75}_{-14.14}$	$103.77(0.31)^{+3.46}_{-5.31}$
3	$40.27(0.05)^{+13.89}_{-9.59}$	$44.55(0.28)^{+1.59}_{-3.08}$	$20.45(0.03)^{+7.09}_{-4.93}$	$23.20(0.16)^{+0.94}_{-1.67}$

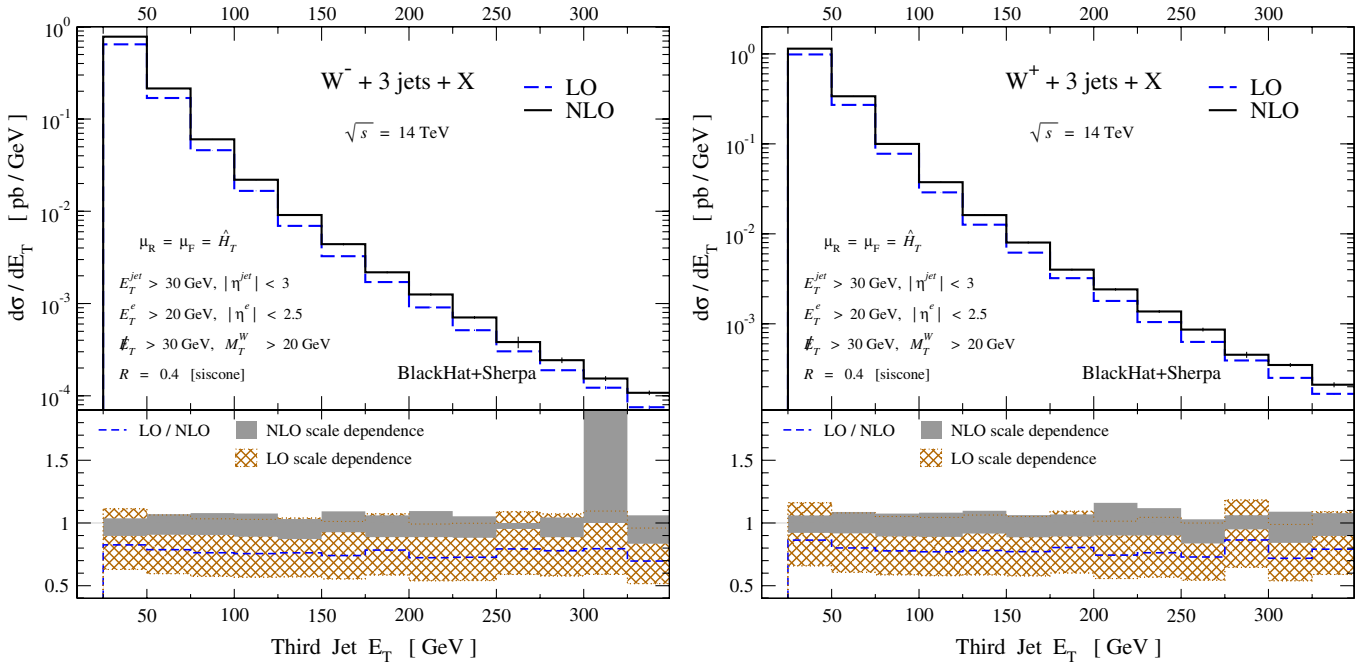

FIG. 17 (color online). The E_T distributions of the third jet, $d\sigma(W \to e\nu + \geq 3\text{-jets})/dE_T^{\rm 3rd\text{-}jet}$, at the LHC. The left panel shows the case of W^- and the right of W^+.

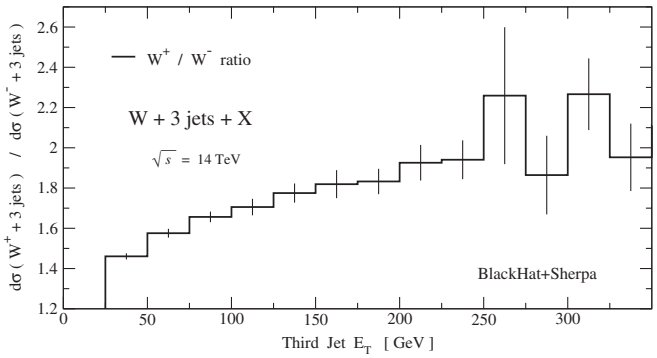

FIG. 18. The ratio of the E_T distribution of the third jet accompanying a W^+ to the same distribution for W^-, evaluated at NLO using $\mu = \hat{H}_T$. The two distributions are shown separately in Fig. 17. The thin vertical lines denote the numerical integration errors.

W^- and W^+ production, except for overall normalization, we generally show a single distribution.

For the inclusive production of $W + 3$ jets, a basic quantity to examine is the E_T distribution for the third-most leading jet in E_T. This distribution is shown in Fig. 17. As in the Tevatron results, the scale uncertainty is considerably reduced at NLO compared to LO. With our default choice of scale $\mu = \hat{H}_T$, the ratio of LO to NLO predictions displayed in the lower panels is rather flat over the entire displayed region. (The upward spike in the NLO band in the W^- plot at 300 GeV is due to a statistical fluctuation in the evaluation at $\mu = \hat{H}_T/2$.) This plot may be compared to the E_T distribution of the second-most energetic jet shown in the right panel of Fig. 9, which undergoes significant shape change between LO and NLO predictions, though less than for the scale choice $\mu = E_T^W$. The dynamic range we show here is larger than in the corresponding plot for the Tevatron.

In order to examine shape differences between the E_T distributions in W^+ and W^- production, in Fig. 18 we show the ratio of the two distributions plotted in Fig. 17. The ratio is greater than unity at low E_T due to the larger total cross section for W^+ production compared to W^-, as given in Tables IV and V. The ratio increases significantly with E_T, on the order of 25% over the range of the plot, because larger E_T forces larger partonic center-of-mass energies, and hence larger values of x where the u quark distribution is more dominant.

The H_T distribution also has slightly different shapes for W^- and W^+ production. The right panel of Fig. 11 shows the H_T distribution in W^- production (with $\mu = \hat{H}_T$). The corresponding plot for W^+ is given in the right panel of Fig. 15. Across the displayed range, the ratio of the NLO W^+ to W^- distributions (not shown) increases slightly. The increase occurs for the same reason as the third jet E_T distribution.

Figure 19 shows the differential distributions with respect to dijet separations ΔR_{ij}. The two hardest jets, labeled 1 and 2, are more likely to be produced in a back-to-back fashion, leading to a more peaked distribution around π. As in other distributions, the NLO scale-dependence band is much smaller than the LO one. The LO and NLO distributions for the separation of the leading two jets are somewhat different from each other in shape. This

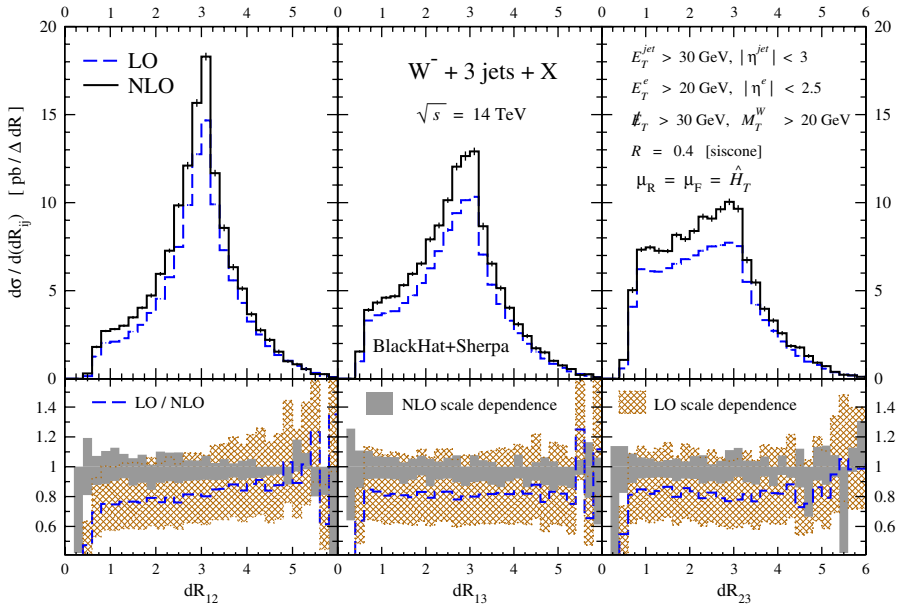

FIG. 19 (color online). The ΔR distributions for W^- + 3-jet production, from left to right: between the first and second, first and third, and second and third jets, using the scale choice $\mu = \hat{H}_T$. The ΔR_{12} distribution between the first and second jets shows a significant shape change in going from LO to NLO, while the other two cases are flat. The distributions for W^+ + 3-jet production are similar.

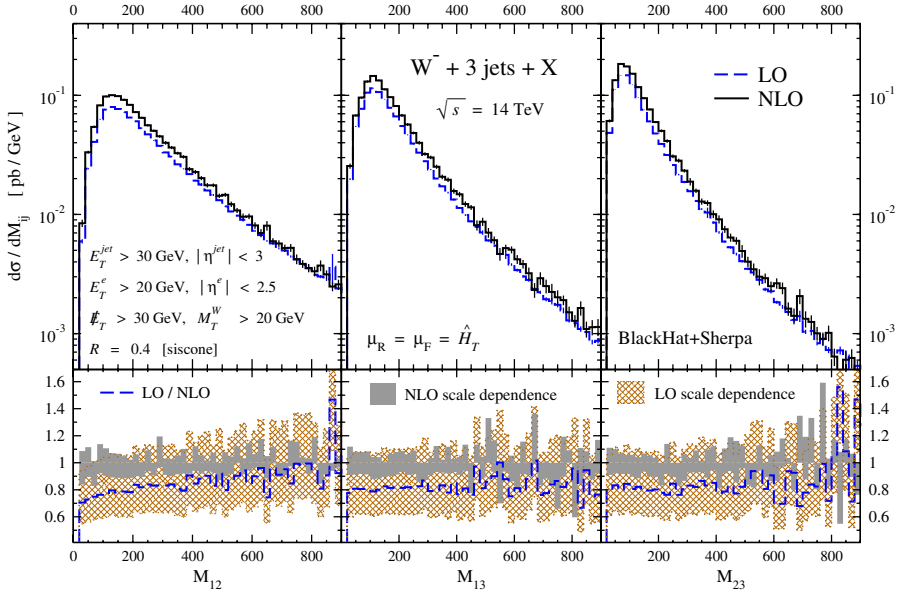

FIG. 20 (color online). The dijet masses in W^- + 3-jet production at the LHC; M_{ij} (in GeV) is the invariant mass of the ith and jth-leading jets, ordered in E_T.

is presumably due to the effect of additional radiation allowing kinematic configurations where the jets are closer together, thereby pushing the weight of the distribution to smaller ΔR values, although the position of the peak is essentially unchanged. The shapes of the other two distributions are similar at LO and NLO. All three distributions show sizable shifts in their overall normalization, for $\mu = \hat{H}_T$.

Figure 20 displays distributions for the dijet masses in W^- + 3-jet production. The three plots in the figure give the dijet mass of the first and second, first and third, and second and third leading jets, denoted by M_{ij} where i and j label the jets. Although our default choice of scale $\mu = \hat{H}_T$ does significantly reduce the shape changes between LO and NLO compared with the choice $\mu = E_T^W$ made for the Tevatron (see Fig. 7), significant shape changes remain for the M_{12} distribution. For the other two cases the ratio between LO and NLO is rather flat. These features have parallels in the ΔR_{ij} distributions in Fig. 19; the physics of the two leading jets is not modeled especially well at LO.

VI. LEPTONS AT THE LHC

We now turn from hadronic observables to leptonic ones. At the LHC, the latter distributions depend strongly on whether a W^+ or a W^- boson has been produced.

Figure 21 shows the pseudorapidity distributions of the daughter charged leptons. Because of the large-x excess of u quarks over d quarks, the qg initial state produces W^+ preferentially, and tends to produce them more forward; this fact accounts for the larger and more forward positron distribution. The lower panels show that in this case, the NLO corrections modify primarily the overall normalization of these distributions, with only a slight change in shape from LO to NLO.

In the right panel of Fig. 16 we showed the positron transverse momentum distribution in W^+ + 3-jet production at the LHC. In order to contrast the distribution with the corresponding distribution in W^- + 3-jet production, Fig. 22 shows the ratio of the NLO transverse energy distributions for the $W^\pm$ boson decay products, charged leptons in the left panel and neutrinos in the right panel. The plots show dramatic differences between the W^+ and W^- distributions, especially for the neutrino E_T, or missing transverse energy. The left panel shows a large ratio for W^+ to W^- at small E_T^e which declines at larger E_T^e. In contrast, the corresponding ratio for the neutrino E_T, or equivalently the missing transverse energy $\not{E}_T$ in the event, starts with a somewhat smaller value but increases rapidly with E_T. This significant increase means that the W^+ + n-jet background to missing-energy-plus-jets signals, when a charged lepton is lost, is more severe than the W^- + n-jet background by a factor of 2 to 3.

This disparate behavior is presumably due to a net left-handed polarization for high E_T $W^\pm$ bosons, which is then analyzed by their leptonic decay via the parity-violating charged-current interaction. Before discussing this situation further, it is useful to recall the dynamics underlying the *longitudinal* (rapidity) charge asymmetry in $W^\pm$ production at the Tevatron, and the corresponding asymmetry for the charged lepton into which the W boson decays [88]. At a $p\bar{p}$ collider, the dominance of u quarks over $\bar{d}$ quarks implies that in the process $u\bar{d} \to W^+ \to e^+ \nu_e$ the W^+ typically moves in the u quark (proton) direction. Because the charged current is left-handed, the u quark must be left-handed, and the $\bar{d}$ antiquark right-handed. In order to

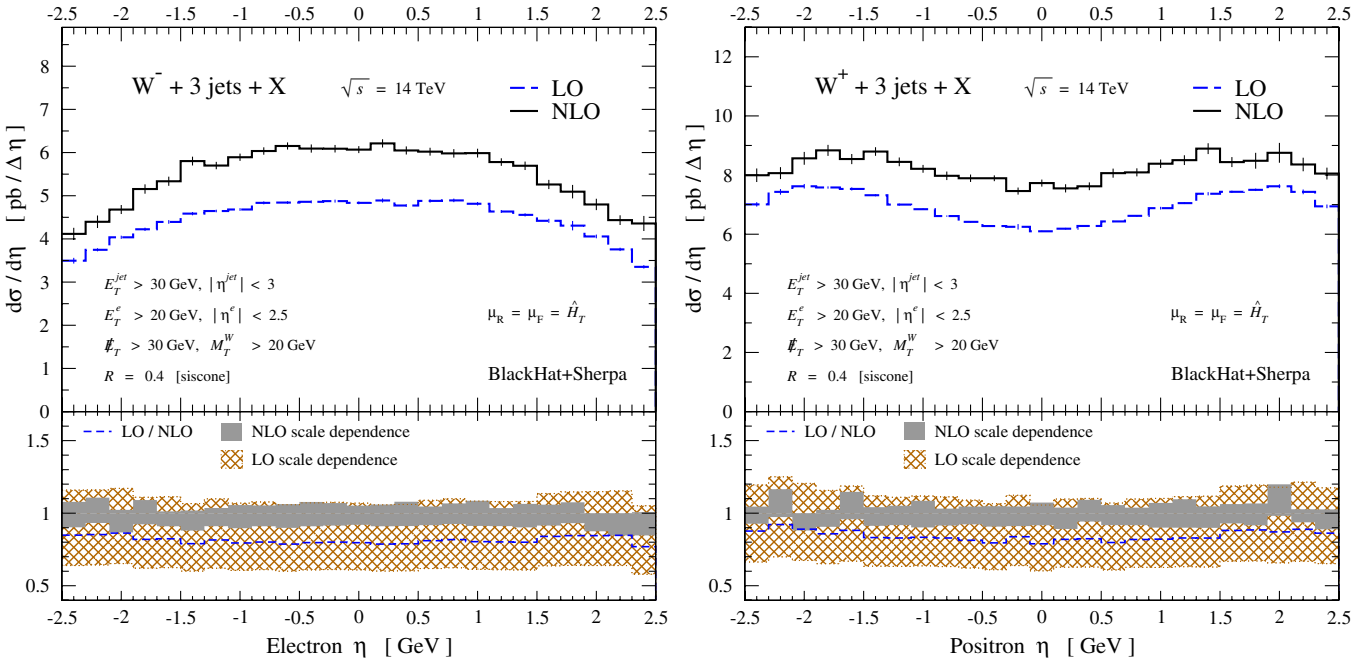

FIG. 21 (color online). The charged-lepton pseudorapidity distribution at the LHC for W^- and W^+ production.

conserve angular momentum, the W^+ must be polarized left-handed along its direction of motion, which in this case (at low transverse energy) is preferentially along the beam axis. In the decay $W^+ \to e^+ \nu_e$, angular-momentum conservation implies that the left-handed W^+ tends to emit the left-handed neutrino forward, and the right-handed positron backward, relative to its direction of motion. The same arguments show that the W^- typically moves in the antiproton direction, is polarized right-handed, and tends to decay with the left-handed electron backward relative to its direction of motion. Both signs of charged leptons are typically more central than are their parent W bosons. In other words, there is a large asymmetry in the rapidity distribution of W^+ bosons at the Tevatron, and a strongly diluted asymmetry in the rapidity distribution of the charged decay lepton [88,89].

Now consider a pp collider (the LHC) and $W^\pm$ bosons moving with large momentum primarily *transverse* to the beam axis, as required to produce the large E_T tails for the decay lepton distributions shown in Fig. 22. Suppose that both W^+ and W^- bosons are polarized left-handed, with a polarization that increases with E_T^W. Then the W^+ will tend to emit the left-handed neutrino forward relative to its direction of motion (resulting in a larger transverse energy) and the right-handed positron backward (smaller transverse energy). In contrast, the W^- will emit the left-handed electron forward. Such decays will produce an enhancement in the neutrino E_T distribution and a depletion in the charged lepton distribution, for W^+ relative to W^-, consistent with the ratios displayed in Fig. 22. We have checked that the same distributions as shown in Fig. 22, but for $W + 1, 2$-jet production, are very similar. Also, the

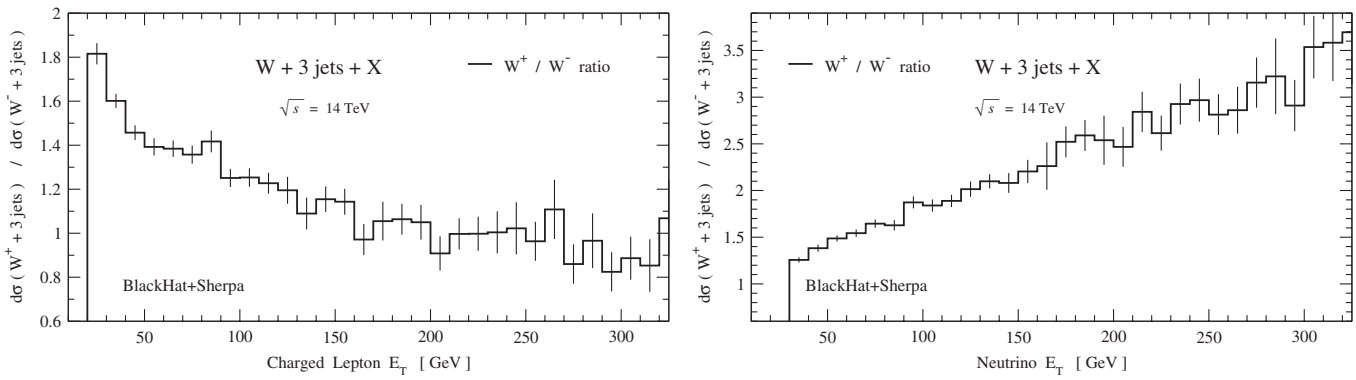

FIG. 22. The left panel shows the ratio of the charged-lepton E_T distributions at the LHC for W^+ and W^- production, evaluated at NLO, while the right panel shows the corresponding ratio for the neutrino E_T, or equivalently the missing transverse energy.

LO ratios in all cases are virtually indistinguishable from the NLO ones. We examined the LO ratios, removing all lepton acceptance cuts, and the same general trends persist (in fact they are even stronger at moderate lepton E_T). The left-handed polarization of both W^+ and W^- is also indicated by the corresponding ratios of W transverse momenta (not shown). These ratios grow monotonically with the W's p_T, but at a lower rate than the ratio for the neutrino E_T. This growth reflects the fact that the larger the W transverse momentum we require, the larger is the required parton momentum fraction x, and thus the more W^+ is favored over W^- by the stiffer u quark distribution.

We do not have a complete understanding of why the W bosons should be polarized left-handed at large transverse momentum, in a manner that is apparently fairly independent of the number of recoiling jets. For W + 1-jet production at LO, it is possible to examine the relevant helicity amplitudes and make such an argument, based on kinematics and on the dominance of the qg initial state. It is also the case that when a very high transverse-momentum (left-handed) quark splits collinearly to another quark and a W boson, with the W boson taking most of the momentum, the W boson is predominantly left-handed. However, these examples certainly do not exhaust all of the possible polarization mechanisms, and a more thorough explanation would require further study.

VII. JET-EMISSION PROBABILITY AT INCREASED PSEUDORAPIDITY SEPARATION

One of the production mechanisms for the Higgs boson at the LHC is via vector-boson fusion [90], which contains partonic subprocesses such as $qQ \to q'Q'H$, mediated by the fusion of two W bosons. Because the Higgs is produced via colorless electroweak vector-boson exchange, a relative absence of radiation is expected between the two forward tagging quark jets, in comparison with QCD background processes with color exchange. A veto on central jets may play an important role in such searches. For this veto to be effective, the background processes should retain a substantial probability of additional radiation, as the pseudorapidity separation $\Delta\eta$ between two of the jets becomes large. The production of a W boson in association with jets is a prime example of a background-type process dominated by color exchange at LO. In Fig. 61 of Ref. [86], a similar question was studied in W + 2-jet and W + 3-jet production at the Tevatron, by looking at the probability of finding a third jet in the acceptance as a function of the pseudorapidity separation of the leading two jets, ordered by transverse energy. In that figure, CDF data was compared with a leading-order QCD prediction.

A more appropriate distribution for assessing the effectiveness of a central jet veto would be to order the jets in pseudorapidity, not E_T, and place an additional constraint that the third jet be between the two most widely separated jets in pseudorapidity, but here we match the choice made in Ref. [86]. In Fig. 23 we plot the ratio of the W^+ + 3-jet to W^+ + 2-jet cross sections at the LHC, as a function of the pseudorapidity separation of the leading two jets, ordered in E_T. This ratio measures the emission probability of a third jet. The solid line gives the NLO equivalent of the LO Tevatron results in Fig. 61 of Ref. [86]. As was found at the Tevatron, the emission probability is substantial, over 20%, and remarkably independent of $\Delta\eta_{12}$. Although we plot only the emission probability for jets accompanying a W^+, the corresponding plot for W^- is essentially indistinguishable from it at NLO. (The difference between the LO and NLO results for W^- is smaller than the difference shown in Fig. 23 for W^+.)

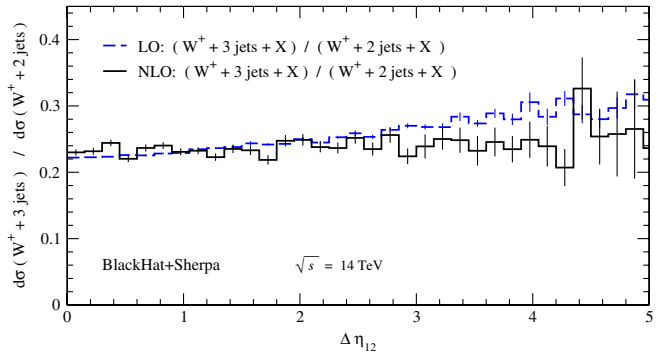

FIG. 23 (color online). The ratio of the inclusive W^+ + 3-jet to W^+ + 2-jet cross sections at NLO, as a function of the pseudorapidity separation $\Delta\eta_{12}$ between the leading two jets at the LHC. The solid (black) line gives the NLO result, while the dashed (blue) line represents the LO result. The ratio for W^- is very similar, particularly at NLO.

VIII. SUBLEADING-COLOR TERMS

In this section, we turn our attention to the question of simplifying a computation by taking advantage of the structure of color sums. As explained in Sec. II, we can organize the matrix elements—leading-order, real-emission, or virtual—in an expansion in $1/N_c$. We expect higher-order terms in this expansion to give smaller contributions numerically; but there are more of them, and their structure is more intricate than that of lower-order terms. In general they take significantly more computer time per event to evaluate. Although one could simply drop these contributions once they have been shown—preferably by direct computation—to be negligible or reliably estimated, we shall describe how to reduce the computer time their direct computation would entail.

In Ref. [34], we used a particular type of "leading-color" approximation (LC NLO), in which a subset of subleading-color terms were dropped. In the real-emission contributions, as well as in the real-subtraction terms, we retained all terms in the color expansion; the same was also true for the singular terms in the virtual matrix elements. The only approximation was within the finite virtual terms.

TABLE VIII. Comparison of LC NLO to full NLO for the total inclusive cross sections in pb of $W + n$-jet production at the Tevatron using CDF's cuts [35] ($E_T^{n\text{th-jet}} > 25$ GeV) and the SISCONE algorithm. For reference (see also Table II), the first column gives the CDF data. The second column shows the LC NLO results and the third column the complete NLO results.

number of jets	CDF	LC NLO	NLO
1	53.5 ± 5.6	$58.3(0.1)^{+4.6}_{-4.6}$	$57.83(0.12)^{+4.36}_{-4.00}$
2	6.8 ± 1.1	$7.81(0.04)^{+0.54}_{-0.91}$	$7.62(0.04)^{+0.62}_{-0.86}$
3	0.84 ± 0.24	$0.908(0.005)^{+0.044}_{-0.142}$	$0.882(0.005)^{+0.057}_{-0.138}$

Here, by "finite" we mean the ϵ^0 term in the Laurent expansion of the infrared-divergent one-loop amplitudes in $\epsilon = (4 - D)/2$, after extracting a multiplicative factor of $c_\Gamma(\epsilon)$ [defined in Eq. (A2)].

The approximation was defined by first dropping the subleading-color terms, that is those suppressed either by powers of $1/N_c^2$ (including those coming from leading-color partial amplitudes) or n_f/N_c (the latter arising from virtual quark loops), in the ratio of the finite virtual terms to the tree-level cross section. In a second step, we multiply this truncated ratio by the tree cross section, with its full color dependence. The net effect of this approximation is to drop quark loops and subleading-color terms in the finite virtual terms that have a different kinematic structure than those at tree level, while retaining the subleading-color terms that have the same kinematic structure. This approximation turns out to be a much better estimate than a strict leading-color approximation (dropping all subleading-color terms in the real-emission terms as well), while still simplifying the calculation considerably by eliminating the need to compute primitive one-loop amplitudes that contribute only to subleading-color terms, such as those shown in Fig. 2.

Table VIII compares NLO results for the total cross sections at the Tevatron with the experimental setup as in Eqs. (3.1) and (3.2) except for the tighter jet cut, $E_T^{\text{jet}} > 25$ GeV. For reference, we also show the corresponding CDF data. The column labeled "LC NLO" contains the results computed using the specific leading-color approximation of Ref. [34]. The last column gives the full NLO result, incorporating all subleading-color terms. Previously [34], we showed explicitly that this approximation is very good for $W + 1$, 2-jet production at the Tevatron, leading to errors of no more than 3%. The entry for $W + 3$ jets is new, and demonstrates that just as for $W + 1$-jet and $W + 2$-jet production, the LC approximation is excellent, shifting the total cross section by just under 3%. This shift is much smaller than the NLO scale dependence. In all cases, the LC NLO and complete NLO result are both in excellent agreement with the data.

Table IX shows results for the total cross section of $W + 3$-jet production at an LHC energy of 14 TeV, using the cuts given in Eqs. (5.1) and (5.2). In this case the column labeled "LC NLO" refers to an LC approximation that is slightly modified from the one used for the Tevatron. We avoid rescaling the leading-color virtual contributions by the ratio of the full-color to leading-color cross section; this allows us to simply add together the "leading-color" and remaining "subleading-color" contributions to obtain

TABLE IX. Comparison of the total cross sections, in pb, between LC NLO and full NLO results for $W + 3$-jet prodiction at the LHC with $\sqrt{s} = 14$ TeV, using the SISCONE algorithm and the cuts of Eqs. (5.1) and (5.2).

cut	W^- LC NLO	W^- NLO	W^+ LC NLO	W^+ NLO
$E_T^{\text{jet}} > 30$ GeV	$28.17(0.13)^{+0.99}_{-2.18}$	$27.52(0.14)^{+1.34}_{-2.81}$	$42.33(0.27)^{+1.82}_{-2.68}$	$41.47(0.27)^{+2.81}_{-3.50}$
$E_T^{\text{jet}} > 40$ GeV	$14.24(0.07)^{+0.76}_{-1.09}$	$13.96(0.07)^{+1.03}_{-1.31}$	$22.08(0.15)^{+1.20}_{-1.44}$	$21.76(0.15)^{+1.68}_{-1.86}$

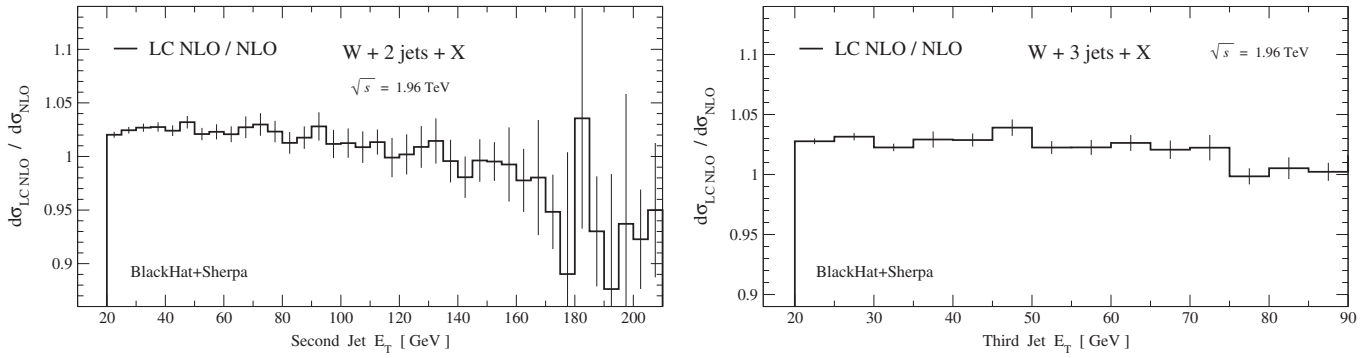

FIG. 24. Ratios of our leading-color approximation to a full-color calculation for E_T distributions in $W + 2$, 3-jet production at the Tevatron. The numerical integration errors on each bin are indicated by thin vertical lines. (Different error sources are combined linearly.)

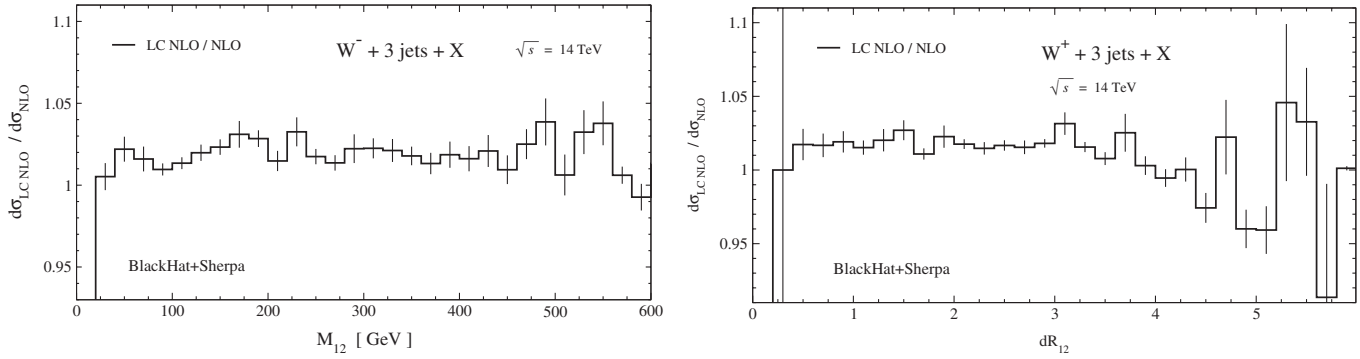

FIG. 25. Ratios of our leading-color approximation to a full-color calculation for the leading-dijet invariant mass distribution in W^- + 3-jet production, and for the ΔR_{12} distribution in W^+ + 3-jet production, at the LHC.

the total cross section. This modification has only a small effect on the corresponding total cross section. Indeed, for Tevatron cross sections the shift is under 1.5% throughout the scale-variation band. (To facilitate comparison to our previously published results [34], for Tevatron cross sections we use the identical LC approximation as in that reference.)

The modified LC approximation is again accurate to 3% for central values, and to 5% for the upper or lower edges of the scale-dependence bands. The good quality of the LC approximation also holds for all distributions we have examined. Examples are shown in Figs. 24 and 25; for both the Tevatron and the LHC, the corrections due to the subleading-color terms are less than 3%, uniformly across the distributions.

The quality of this approximation has important implications for organizing the calculation of virtual contributions. As discussed in section II B, the primitive amplitudes entering the leading-color approximation used in Ref. [34] and discussed above are a small subset of the primitive amplitudes required for the complete virtual correction. The full result, including subleading-color terms, requires the computation of a much larger number of primitive amplitudes, 336 for two-quark processes and 80 for four-quark processes[7] as opposed to 48 for two-quark processes and 8 for four-quark processes in the leading-color computation. Furthermore, the subleading-color primitive amplitudes shown in Fig. 2 have a more complicated analytic structure than the leading-color ones shown in Fig. 1, because the W does not have to be ordered with respect to parton(s) emitted between the W boson and the q or q'. The partial lack of color-ordering implies cuts and poles in more channels. As a result, it takes about 27 times longer per phase-space point to evaluate the subleading-color virtual terms than the leading-color ones.

This large factor may seem to be a cause for concern. However, the smallness of the subleading-color contributions, discussed above, comes to our rescue. If we require that the numerical integration errors due to the subleading-color contributions be comparable to those coming from the leading-color ones, we can allow for larger *relative* errors in the evaluation of the subleading-color terms. We can adopt a ''color-expansion sampling'' approach wherein we can use far fewer phase-space points (typically a factor of 20 fewer) to evaluate them [49] as compared to the leading-color terms. (We must ensure that there are sufficient statistics in each bin of every distribution of interest, of course.) There is no need to know ahead of time what the relative size of the two contributions is; we simply stop the integration when the desired numerical precision is reached for each contribution separately. This approach requires only a bit more than a factor of 2 more computer time for the full-color result than for the leading-color approximation, with our present setup. It saves roughly a factor of 30 in computer time, compared to the naive approach of evaluating the subleading-color terms at every phase-space point. We expect to obtain further improvements in the evaluation efficiency of subleading-color contributions through improved reuse of primitive amplitudes. Together with color-expansion sampling, this should reduce the time required for computing the subleading-color terms to a small fraction of the total computer time. This color-expansion sampling approach would naturally be implemented in a dynamical way by treating the subleading-color contributions as another set of ''subprocesses'' within the SHERPA multichannel integration.

Were we seeking to optimize the computer time more aggressively, we would let the total error be *dominated* by the most time-consuming part of the calculation, namely, the subleading-color terms. We have, however, opted for a more conservative approach, keeping the integration error from the subleading-color contributions in line with the errors from the other contributions.

IX. CONCLUSIONS

In this paper, we presented the first complete NLO study of W + 3-jet production, incorporating all massless par-

[7]The count for the four-quark case is for nonidentical fermions; for identical fermions, multiply by 2.

tonic contributions. We compared the total cross section and the third-jet E_T distribution to Tevatron data [35]. We also presented NLO predictions for the H_T and dijet mass distributions. It will be interesting to compare these and other distributions to Tevatron data sets with larger statistics than those already published by CDF [35]. We presented a variety of distributions at the declared final running energy of the LHC, including many relevant for standard model backgrounds to events with large missing energy and to Higgs boson production via vector-boson fusion. As expected, we find a much smaller renormalization- and factorization-scale dependence in all distributions at NLO, compared to LO results. Although the LHC will start running at lower energy, our choice should help facilitate comparisons to earlier studies based on leading-order QCD and matching to parton showers [2–4,7,8,84,91].

We have shown explicitly in $W + 3$ jet distributions that the scale dependence of LO predictions is not restricted to overall normalizations. An infelicitous choice of scale can change the shapes of distributions substantially between LO and NLO. This effect is much more pronounced at the LHC than at the Tevatron. One can reduce the change in shape of distributions between LO and NLO by choosing a scale dynamically, event by event, corresponding to a typical scale for the event, as noted in, for example, Refs. [14,16,38]. The problem with poor scale choices can be much more severe than just changes in shape between LO and NLO results. Indeed, for sufficiently poor choices, such as the fixed scale $\mu = M_W$ or the transverse energy of the W boson, $\mu = E_T^W$, large logarithms can appear in some distributions, invalidating even an NLO prediction. We find that the total (partonic) transverse energy $\hat{H}_T$ is a more appropriate scale choice for $W + 3$-jet production than the W transverse energy or the fixed scales used in previous Tevatron analyses. (A fixed fraction of the total transverse energy would also be appropriate.) We expect that this scale choice will be appropriate to a variety of higher-multiplicity processes, and recommend its use in LO predictions (when an NLO one is not available) as well as at NLO. A recent paper [40] motivates a similar type of scale choice using soft-collinear effective theory, and we have contrasted its properties with those of $\hat{H}_T$. Of course, a simple scale choice is no substitute for a complete NLO prediction. In some distributions, such as the transverse energy of the second-most energetic jet and the ΔR separation between the two leading jets, the NLO calculation incorporates physics effects that are not captured by simple changes of scale.

We also confirmed that our previous NLO analysis of $W + 3$-jet production [34], which used a specific leading-color approximation, is valid to within 3%. This error is quite a bit smaller than other uncertainties, such as that implicit in the scale dependence, or that due to uncertainties in the parton distribution functions. However, we can draw this conclusion only after computing the subleading-color terms, as we have done here. To evaluate the subleading-color terms efficiently, we used "color-expansion sampling." The subleading-color terms require much more computer time per phase space point. However, because they are small, only a few percent of the leading-color ones, we can tolerate a much larger relative error for them from the Monte Carlo integration, thus sampling them much less often. We expect that this general approach will be an effective technique for reducing the computer-time requirements for ever-more complicated processes such as $Z + 3$-jet, $W + 4$-jet, or $Z + 4$-jet production.

In our analyses we mainly used the SISCONE jet algorithm; we also presented total cross sections using the k_T jet algorithm [87] at the LHC. These jet algorithms are infrared-safe to all orders in perturbation theory. With our setup it is a simple matter to replace one infrared-safe cone algorithm with any other desired one. We defer a study of the anti-k_T algorithm [92], which has certain experimental advantages such as uniform catchment areas for soft radiation, to future work. From a perturbative viewpoint, infrared safety is essential; infrared-unsafe quantities are simply logarithmically divergent. In the real world, perturbation theory does not go on forever but is overtaken by nonperturbative dynamics around the confinement scale. Infrared-unsafe quantities are not infinite, but the infinities are cut off and replaced by quantities determined by nonperturbative physics. The logarithms translate [37,93] into inverse powers of the strong coupling α_S, thereby spoiling the perturbative expansion. This is an important practical problem because the jet algorithms traditionally used at the Tevatron by the CDF and D0 collaborations are, in fact, infrared-unsafe beyond the lowest orders [37]. Unknown nonperturbative corrections for these algorithms would undo many of the benefits of a higher-order prediction, especially in the context of new, higher statistics data. Accordingly, it is highly desirable that future experimental analyses at both the Tevatron and the LHC use an infrared-safe jet algorithm.

Our paper demonstrates the utility of on-shell methods for computing one-loop matrix elements entering state-of-the-art NLO QCD predictions for processes of phenomenological interest at the LHC. We used BLACKHAT, an efficient new code library based on these methods. The NLO $W + 3$-jet results reported here also demonstrate the functionality of our computational setup, which uses BLACKHAT in conjunction with the SHERPA package. Besides handling the real-emission contributions and infrared-singular phase space via AMEGIC++, the SHERPA framework offers a convenient set of tools for integrating over phase space and analyzing the results.

There are many relevant processes with large numbers of final-state objects such as jets that remain to be computed, especially those involving vector bosons, jets, heavy

quarks and Higgs bosons [9]. Such processes are backgrounds to the production of new heavy particles with multibody decays. Our setup is robust enough to deal systematically with such processes. In the present paper, we have demonstrated the new tools and on-shell methods at work for the nontrivial case of W + 3-jet production at hadron colliders. We look forward to comparing our predictions against forthcoming LHC data.

ACKNOWLEDGMENTS

We thank Jeppe Andersen, Christian Bauer, John Campbell, Keith Ellis, Beate Heinemann, Joey Huston, Pavel Nadolsky, Michael Peskin, Gavin Salam, Rainer Wallny, and Giulia Zanderighi for helpful discussions. We especially thank Costas Papadopoulos and Roberto Pittau for assistance in comparing results for the virtual contributions to squared matrix elements. This research was supported by the U.S. Department of Energy under Contract Nos. DE-FG03-91ER40662, DE-AC02-76SF00515, and DE-FC02-94ER40818. DAK's research is supported by the European Research Council under Advanced Investigator Grant No. ERC-AdG-228301. HI's work is supported by the U.S. LHC Theory Initiative through NSF Grant No. PHY-0705682. This research used resources of Academic Technology Services at UCLA, PhenoGrid using the GridPP infrastructure, and the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

APPENDIX: SQUARED MATRIX ELEMENTS AT ONE POINT IN PHASE SPACE

In order to aid future implementations of virtual corrections for W + 3-jet production in other numerical codes, we present values of the one-loop virtual corrections to the squared matrix elements, $d\sigma_V^{(1)}$, at one point in phase space. This comes from the interference between the tree and one-loop amplitudes, summed over all colors and helicities, for $N_c = 3$ and $n_f = 5$ massless quark flavors.

In Table X we present numerical values for four representative subprocesses. All other subprocesses are related to these four by crossing symmetry. In the second and third lines of Table X, the presence of two identical quarks (after crossing all particles into the final state) means that amplitudes are antisymmetrized under exchange of the two.

We quote numerical results for the ultraviolet-renormalized virtual corrections in the 't Hooft-Veltman variant of dimensional regularization [94]. The remaining singularities in the dimensional regularization parameter $\epsilon = (4 - D)/2$ arise from the virtual soft and collinear singularities in the one-loop amplitudes.

The quoted values are for the ratio of the virtual corrections to the tree-level squared matrix element $d\sigma^{(0)}$. Explicitly, we define the ratio,

$$\widehat{d\sigma}_V^{(1)} \equiv \frac{1}{8\pi\alpha_S c_\Gamma(\epsilon)} \frac{d\sigma_V^{(1)}}{d\sigma^{(0)}}, \tag{A1}$$

where we have also separated out the dependence on the strong coupling α_S and the overall factor $c_\Gamma(\epsilon)$, defined by

$$c_\Gamma(\epsilon) = \frac{1}{(4\pi)^{2-\epsilon}} \frac{\Gamma(1+\epsilon)\Gamma^2(1-\epsilon)}{\Gamma(1-2\epsilon)}. \tag{A2}$$

The coupling constants, mass and width of the W boson are given in Sec. II G. However, the numerical values for the ratio (A1), given in Table X, are independent of these parameters; coupling constants as well as the W boson Breit-Wigner factor cancel between the tree and virtual correction terms.

We choose the phase-space point given in Eqs. (9.3) and (9.4) of Ref. [61],

$$\begin{aligned}
k_1 &= \frac{\mu}{2}(1, -\sin\theta, -\cos\theta\sin\phi, -\cos\theta\cos\phi),\\
k_2 &= \frac{\mu}{2}(1, \sin\theta, \cos\theta\sin\phi, \cos\theta\cos\phi),\\
k_3 &= \frac{\mu}{3}(1, 1, 0, 0), \qquad k_4 = \frac{\mu}{8}(1, \cos\beta, \sin\beta, 0),\\
k_5 &= \frac{\mu}{10}(1, \cos\alpha\cos\beta, \cos\alpha\sin\beta, \sin\alpha),\\
k_6 &= \frac{\mu}{12}(1, \cos\gamma\cos\beta, \cos\gamma\sin\beta, \sin\gamma),\\
k_7 &= k_1 + k_2 - k_3 - k_4 - k_5 - k_6,
\end{aligned}$$

where

TABLE X. Numerical values of the normalized virtual correction to the squared matrix elements, $\widehat{d\sigma}_V^{(1)}$, at the phase-space point given in the text, for the four basic partonic subprocesses for W + 3-jet production at a hadron collider. We give the finite parts along with the coefficients of the poles in ϵ.

$\widehat{d\sigma}_V^{(1)}$	$1/\epsilon^2$	$1/\epsilon$	ϵ^0
$(1_{\bar{u}}2_c \to 3_c4_{\bar{d}}5_g6_{e^-}7_{\bar{\nu}})$	−8.333 333 333	−32.376 772 10	1.778 061 330
$(1_{\bar{u}}2_u \to 3_u4_{\bar{d}}5_g6_{e^-}7_{\bar{\nu}})$	−8.333 333 333	−32.408 071 65	1.035 000 256
$(1_{\bar{u}}2_d \to 3_d4_{\bar{d}}5_g6_{e^-}7_{\bar{\nu}})$	−8.333 333 333	−32.507 501 36	0.478 803 062 4
$(1_{\bar{u}}2_g \to 3_g4_g5_{\bar{d}}6_{e^-}7_{\bar{\nu}})$	−11.666 666 67	−42.343 036 28	−13.979 912 25

$$\theta = \frac{\pi}{4}, \qquad \phi = \frac{\pi}{6}, \qquad \alpha = \frac{\pi}{3},$$
$$\gamma = \frac{2\pi}{3}, \qquad \cos\beta = -\frac{37}{128},$$

and the renormalization scale μ is set to $\mu = 7$ GeV. We have flipped the signs of k_1 and k_2 compared to Ref. [61], to correspond to $2 \to 5$ kinematics, instead of $0 \to 7$ kinematics. The labeling of the parton and lepton momenta is indicated explicitly in the first column of Table X.

[1] S. D. Ellis, R. Kleiss, and W. J. Stirling, Phys. Lett. **158B**, 341 (1985); M. L. Mangano, Eur. Phys. J. C **59**, 373 (2009).
[2] T. Stelzer and W. F. Long, Comput. Phys. Commun. **81**, 357 (1994); A. Pukhov *et al.*, arXiv:hep-ph/9908288; M. L. Mangano, M. Moretti, F. Piccinini, R. Pittau, and A. D. Polosa, J. High Energy Phys. 07 (2003) 001.
[3] A. Kanaki and C. G. Papadopoulos, Comput. Phys. Commun. **132**, 306 (2000); A. Cafarella, C. G. Papadopoulos, and M. Worek, arXiv:0710.2427.
[4] F. Krauss, R. Kuhn, and G. Soff, J. High Energy Phys. 02 (2002) 044.
[5] H. U. Bengtsson and T. Sjöstrand, Comput. Phys. Commun. **46**, 43 (1987); T. Sjöstrand, P. Eden, C. Friberg, L. Lönnblad, G. Miu, S. Mrenna, and E. Norrbin, Comput. Phys. Commun. **135**, 238 (2001); G. Marchesini and B. R. Webber, Cavendish Report No. HEP-87/9; G. Marchesini, B. R. Webber, G. Abbiendi, I. G. Knowles, M. H. Seymour, and L. Stanco, Comput. Phys. Commun. **67**, 465 (1992); G. Corcella *et al.*, arXiv:hep-ph/0210213.
[6] T. Gleisberg, S. Höche, F. Krauss, M. Schönherr, S. Schumann, F. Siegert, and J. Winter, J. High Energy Phys. 02 (2009) 007.
[7] S. Catani, F. Krauss, R. Kuhn, and B. R. Webber, J. High Energy Phys. 11 (2001) 063; M. Mangano, in Fermilab ME/MC Tuning Workshop, 2004 (to be published).
[8] M. L. Mangano, M. Moretti, F. Piccinini, and M. Treccani, J. High Energy Phys. 01 (2007) 013; S. Mrenna and P. Richardson, J. High Energy Phys. 05 (2004) 040.
[9] Z. Bern *et al.*, arXiv:0803.0494.
[10] S. Frixione and B. R. Webber, J. High Energy Phys. 06 (2002) 029; S. Frixione, P. Nason, and B. R. Webber, J. High Energy Phys. 08 (2003) 007; S. Frixione, P. Nason, and C. Oleari, J. High Energy Phys. 11 (2007) 070; S. Alioli, P. Nason, C. Oleari, and E. Re, J. High Energy Phys. 07 (2008) 060.
[11] H. Fritzsch and P. Minkowski, Phys. Lett. **73B**, 80 (1978); G. Altarelli, G. Parisi, and R. Petronzio, Phys. Lett. **76B**, 351 (1978); K. Kajantie and R. Raitio, Nucl. Phys. **B139**, 72 (1978); K. Kajantie, J. Lindfors, and R. Raitio, Nucl. Phys. **B144**, 422 (1978).
[12] R. Kleiss and W. J. Stirling, Nucl. Phys. **B262**, 235 (1985).
[13] S. D. Ellis, R. Kleiss, and W. J. Stirling, Phys. Lett. **154B**, 435 (1985).
[14] M. L. Mangano and S. J. Parke, Phys. Rev. D **41**, 59 (1990).
[15] R. K. Ellis, G. Martinelli, and R. Petronzio, Nucl. Phys. **B211**, 106 (1983); W. T. Giele, E. W. N. Glover, and D. A. Kosower, Nucl. Phys. **B403**, 633 (1993).
[16] P. B. Arnold and M. H. Reno, Nucl. Phys. **B319**, 37 (1989); **B330**, 284(E) (1990).
[17] F. A. Berends, H. Kuijf, B. Tausk, and W. T. Giele, Nucl. Phys. **B357**, 32 (1991).
[18] F. Abe *et al.* (CDF Collaboration), Phys. Rev. Lett. **74**, 2626 (1995); S. Abachi *et al.* (D0 Collaboration), Phys. Rev. Lett. **74**, 2632 (1995).
[19] Z. Bern, L. J. Dixon, and D. A. Kosower, Nucl. Phys. **B513**, 3 (1998).
[20] Z. Bern, L. J. Dixon, D. C. Dunbar, and D. A. Kosower, Nucl. Phys. **B425**, 217 (1994); Nucl. Phys. **B435**, 59 (1995).
[21] E. W. N. Glover and D. J. Miller, Phys. Lett. B **396**, 257 (1997); Z. Bern, L. J. Dixon, D. A. Kosower, and S. Weinzierl, Nucl. Phys. **B489**, 3 (1997); J. M. Campbell, E. W. N. Glover, and D. J. Miller, Phys. Lett. B **409**, 503 (1997).
[22] J. M. Campbell and R. K. Ellis, Phys. Rev. D **65**, 113007 (2002).
[23] W. T. Giele, S. Keller, and E. Laenen, Phys. Lett. B **372**, 141 (1996).
[24] J. M. Campbell, R. K. Ellis, F. Maltoni, and S. Willenbrock, Phys. Rev. D **75**, 054015 (2007).
[25] F. Febres Cordero, L. Reina, and D. Wackeroth, Phys. Rev. D **74**, 034007 (2006); **80**, 034015 (2009).
[26] J. M. Campbell, R. K. Ellis, F. Febres Cordero, F. Maltoni, L. Reina, D. Wackeroth, and S. Willenbrock, Phys. Rev. D **79**, 034023 (2009).
[27] Z. Bern, L. J. Dixon, and D. A. Kosower, Ann. Phys. (N.Y.) **322**, 1587 (2007).
[28] A. Bredenstein, A. Denner, S. Dittmaier, and S. Pozzorini, J. High Energy Phys. 08 (2008) 108; Phys. Rev. Lett. **103**, 012002 (2009).
[29] G. Ossola, C. G. Papadopoulos, and R. Pittau, J. High Energy Phys. 03 (2008) 042.
[30] C. F. Berger, Z. Bern, L. J. Dixon, F. Febres Cordero, D. Forde, H. Ita, D. A. Kosower, and D. Maître, Phys. Rev. D **78**, 036003 (2008).
[31] W. T. Giele and G. Zanderighi, J. High Energy Phys. 06 (2008) 038.
[32] A. Lazopoulos, arXiv:0812.2998; J. C. Winter and W. T. Giele, arXiv:0902.0094.
[33] A. van Hameren, C. G. Papadopoulos, and R. Pittau, arXiv:0903.4665.
[34] C. F. Berger *et al.*, Phys. Rev. Lett. **102**, 222001 (2009).
[35] T. Aaltonen *et al.* (CDF Collaboration), Phys. Rev. D **77**, 011108 (2008).
[36] F. Abe *et al.* (CDF Collaboration), Phys. Rev. Lett. **75**,

3997 (1995).
[37] G. P. Salam and G. Soyez, J. High Energy Phys. 05 (2007) 086.
[38] S. Frixione, Nucl. Phys. **B410**, 280 (1993); U. Baur, T. Han, and J. Ohnemus, Phys. Rev. D **53**, 1098 (1996); **57**, 2823 (1998); L. J. Dixon, Z. Kunszt, and A. Signer, Phys. Rev. D **60**, 114037 (1999); G. Bozzi, B. Jäger, C. Oleari, and D. Zeppenfeld, Phys. Rev. D **75**, 073004 (2007).
[39] F. Abe *et al.* (CDF Collaboration), Phys. Rev. Lett. **77**, 448 (1996); T. Aaltonen *et al.* (CDF-Run II Collaboration), Phys. Rev. Lett. **100**, 102001 (2008).
[40] C. W. Bauer and B. O. Lange, arXiv:0905.4739.
[41] R. K. Ellis, K. Melnikov, and G. Zanderighi, J. High Energy Phys. 04 (2009) 077.
[42] R. K. Ellis, K. Melnikov, and G. Zanderighi, arXiv:0906.1445.
[43] R. K. Ellis, W. T. Giele, Z. Kunszt, and K. Melnikov, arXiv:0806.3467.
[44] W. T. Giele, Z. Kunszt, and K. Melnikov, J. High Energy Phys. 04 (2008) 049.
[45] C. F. Berger, Z. Bern, L. J. Dixon, F. Febres Cordero, D. Forde, H. Ita, D. A. Kosower, and D. Maître, arXiv:0808.0941.
[46] S. Catani and M. H. Seymour, Phys. Lett. B **378**, 287 (1996); Nucl. Phys. **B485**, 291 (1997); **B510**, 503(E) (1998).
[47] T. Gleisberg and F. Krauss, Eur. Phys. J. C **53**, 501 (2008).
[48] M. H. Seymour and C. Tevlin, arXiv:0803.2231; K. Hasegawa, S. Moch, and P. Uwer, Nucl. Phys. B, Proc. Suppl. **183**, 268 (2008); R. Frederix, T. Gehrmann, and N. Greiner, J. High Energy Phys. 09 (2008) 122; M. Czakon, C. G. Papadopoulos, and M. Worek, J. High Energy Phys. 08 (2009) 085.
[49] L. J. Dixon and A. Signer, Phys. Rev. D **56**, 4031 (1997).
[50] R. Kleiss and R. Pittau, Comput. Phys. Commun. **83**, 141 (1994).
[51] Z. Bern, L. J. Dixon, and D. A. Kosower, Nucl. Phys. **B437**, 259 (1995).
[52] R. Britto, F. Cachazo, and B. Feng, Nucl. Phys. **B715**, 499 (2005); R. Britto, F. Cachazo, B. Feng, and E. Witten, Phys. Rev. Lett. **94**, 181602 (2005).
[53] Z. Bern, L. J. Dixon, and D. A. Kosower, Phys. Rev. D **71**, 105013 (2005); **72**, 125003 (2005); **73**, 065013 (2006); D. Forde and D. A. Kosower, Phys. Rev. D **73**, 065007 (2006); **73**, 061701 (2006); C. F. Berger, Z. Bern, L. J. Dixon, D. Forde, and D. A. Kosower, Phys. Rev. D **75**, 016006 (2007).
[54] R. Britto, F. Cachazo, and B. Feng, Nucl. Phys. **B725**, 275 (2005).
[55] G. Ossola, C. G. Papadopoulos, and R. Pittau, Nucl. Phys. **B763**, 147 (2007).
[56] R. K. Ellis, W. T. Giele, and Z. Kunszt, J. High Energy Phys. 03 (2008) 003.
[57] D. Forde, Phys. Rev. D **75**, 125019 (2007).
[58] R. J. Eden, P. V. Landshoff, D. I. Olive, and J. C. Polkinghorne, *The Analytic S Matrix* (Cambridge University Press, Cambridge, England, 1966).
[59] W. T. Giele and E. W. N. Glover, J. High Energy Phys. 04 (2004) 029; T. Binoth, J. P. Guillet, G. Heinrich, E. Pilon, and C. Schubert, J. High Energy Phys. 10 (2005) 015; A. Denner and S. Dittmaier, Nucl. Phys. **B734**, 62 (2006); T. Binoth, J. P. Guillet, G. Heinrich, E. Pilon, and T. Reiter, arXiv:0810.0992.
[60] T. Binoth *et al.*, arXiv:0807.0605.
[61] C. F. Berger, Z. Bern, L. J. Dixon, D. Forde, and D. A. Kosower, Phys. Rev. D **74**, 036009 (2006).
[62] Z. Bern, L. J. Dixon, and D. A. Kosower, Annu. Rev. Nucl. Part. Sci. **46**, 109 (1996).
[63] L. M. Brown and R. P. Feynman, Phys. Rev. **85**, 231 (1952); L. M. Brown, Nuovo Cimento **21**, 3878 (1961); B. Petersson, J. Math. Phys. (N.Y.) **6**, 1955 (1965); G. Källén and J. S. Toll, J. Math. Phys. (N.Y.) **6**, 299 (1965); D. B. Melrose, Nuovo Cimento **40**, 181 (1965); G. Passarino and M. J. G. Veltman, Nucl. Phys. **B160**, 151 (1979); W. L. van Neerven and J. A. M. Vermaseren, Phys. Lett. **137B**, 241 (1984); Z. Bern, L. J. Dixon, and D. A. Kosower, Phys. Lett. B **302**, 299 (1993); **318**, 649(E) (1993); J. Fleischer, F. Jegerlehner, and O. V. Tarasov, Nucl. Phys. **B566**, 423 (2000); T. Binoth, J. P. Guillet, and G. Heinrich, Nucl. Phys. **B572**, 361 (2000); G. Duplančić and B. Nižić, Eur. Phys. J. C **35**, 105 (2004).
[64] G. 't Hooft and M. J. G. Veltman, Nucl. Phys. **B153**, 365 (1979); G. J. van Oldenborgh and J. A. M. Vermaseren, Z. Phys. C **46**, 425 (1990); W. Beenakker and A. Denner, Nucl. Phys. **B338**, 349 (1990); A. Denner, U. Nierste, and R. Scharf, Nucl. Phys. **B367**, 637 (1991); Z. Bern, L. J. Dixon, and D. A. Kosower, Nucl. Phys. **B412**, 751 (1994); T. Hahn and M. Pérez-Victoria, Comput. Phys. Commun. **118**, 153 (1999); R. K. Ellis and G. Zanderighi, J. High Energy Phys. 02 (2008) 002.
[65] F. A. Berends, R. Kleiss, P. De Causmaecker, R. Gastmans, and T. T. Wu, Phys. Lett. **103B**, 124 (1981); P. De Causmaecker, R. Gastmans, W. Troost, and T. T. Wu, Nucl. Phys. **B206**, 53 (1982); Z. Xu, D. H. Zhang, and L. Chang, TUTP-84/3-TSINGHUA; J. F. Gunion and Z. Kunszt, Phys. Lett. **161B**, 333 (1985); Z. Xu, D. H. Zhang, and L. Chang, Nucl. Phys. **B291**, 392 (1987).
[66] E. Witten, Commun. Math. Phys. **252**, 189 (2004).
[67] Z. Bern and A. G. Morgan, Nucl. Phys. **B467**, 479 (1996); Z. Bern, L. J. Dixon, D. C. Dunbar, and D. A. Kosower, Phys. Lett. B **394**, 105 (1997).
[68] C. Anastasiou, R. Britto, B. Feng, Z. Kunszt, and P. Mastrolia, Phys. Lett. B **645**, 213 (2007); R. Britto and B. Feng, J. High Energy Phys. 02 (2008) 095.
[69] G. Ossola, C. G. Papadopoulos, and R. Pittau, J. High Energy Phys. 05 (2008) 004.
[70] W. L. van Neerven, Nucl. Phys. **B268**, 453 (1986).
[71] R. K. Ellis, W. T. Giele, Z. Kunszt, K. Melnikov, and G. Zanderighi, J. High Energy Phys. 01 (2009) 012.
[72] S. D. Badger, J. High Energy Phys. 01 (2009) 049.
[73] M. L. Mangano and S. J. Parke, Phys. Rep. **200**, 301 (1991); L. J. Dixon, in *QCD & Beyond: Proceedings of TASI '95*, edited by D. E. Soper (World Scientific, Singapore, 1996).
[74] W. B. Kilgore, arXiv:0711.5015; R. Britto, B. Feng, and P. Mastrolia, Phys. Rev. D **78**, 025031 (2008); S. D. Badger, Nucl. Phys. B, Proc. Suppl. **183**, 220 (2008); R. Britto and B. Feng, arXiv:0904.2766.
[75] Y. Hida, X. S. Li, and D. H. Bailey, Report No. LBNL-46996, http://crd.lbl.gov/~dhbailey/mpdist.
[76] M. Seymour (private communication).
[77] Z. Nagy, Phys. Rev. D **68**, 094002 (2003).

[78] G. P. Lepage, J. Comput. Phys. **27**, 192 (1978).
[79] J. Pumplin *et al.*, J. High Energy Phys. 07 (2002) 012.
[80] W. T. Giele and E. W. N. Glover, Phys. Rev. D **46**, 1980 (1992); Z. Kunszt, A. Signer and Z. Trócsányi, Nucl. Phys. **B420**, 550 (1994).
[81] B. D. Cooper, Ph.D. thesis, FERMILAB-THESIS-2006-61.
[82] F. Abe *et al.* (CDF Collaboration), Phys. Rev. D **45**, 1448 (1992).
[83] R. Akers *et al.* (OPAL Collaboration), Z. Phys. C **63**, 197 (1994); M. H. Seymour, Nucl. Phys. **B513**, 269 (1998).
[84] S. Höche, F. Krauss, N. Lavesson, L. Lönnblad, M. Mangano, A. Schälicke, and S. Schumann, arXiv:hep-ph/0602031; J. Alwall *et al.*, Eur. Phys. J. C **53**, 473 (2008).
[85] D. Stump, J. Huston, J. Pumplin, W. K. Tung, H. L. Lai, S. Kuhlmann, and J. F. Owens, J. High Energy Phys. 10 (2003) 046.
[86] J. M. Campbell, J. W. Huston, and W. J. Stirling, Rep. Prog. Phys. **70**, 89 (2007).
[87] S. Catani, Y. L. Dokshitzer, M. H. Seymour, and B. R. Webber, Nucl. Phys. **B406**, 187 (1993); S. D. Ellis and D. E. Soper, Phys. Rev. D **48**, 3160 (1993).
[88] R. K. Ellis, W. J. Stirling, and B. R. Webber, *QCD and Collider Physics* (Cambridge University Press, Cambridge, England, 1996).
[89] T. Aaltonen *et al.* (CDF Collaboration), Phys. Rev. Lett. **102**, 181801 (2009).
[90] D. Zeppenfeld, R. Kinnunen, A. Nikitenko, and E. Richter-Was, Phys. Rev. D **62**, 013009 (2000); M. Dührssen, S. Heinemeyer, H. Logan, D. Rainwater, G. Weiglein, and D. Zeppenfeld, Phys. Rev. D **70**, 113009 (2004).
[91] F. Krauss, A. Schälicke, S. Schumann, and G. Soff, Phys. Rev. D **70**, 114009 (2004); **72**, 054017 (2005); N. Lavesson and L. Lönnblad, J. High Energy Phys. 07 (2005) 054.
[92] M. Cacciari, G. P. Salam, and G. Soyez, J. High Energy Phys. 04 (2008) 063.
[93] G. P. Salam, arXiv:0906.1833.
[94] G. 't Hooft and M. Veltman, Nucl. Phys. **B44**, 189 (1972).